ADSTAR Distributed Storage Manager
for Windows NT**

IBM

Quick Start

Version 2

```
┌─ Note! ──────────────────────────────────────────────────────────────────┐
│                                                                            │
│  Before using this information and the product it supports, be sure to read the general information under │
│  Appendix D, "Notices" on page 55.                                         │
│                                                                            │
└────────────────────────────────────────────────────────────────────────────┘
```

This book is also available in a softcopy form that can be viewed with the IBM BookManager READ licensed program.

Contents

Before You Begin

ADSTAR Distributed Storage Manager (ADSM) is a client/server program that provides storage management services to customers in a multivendor computer environment. ADSM provides an automated, centrally scheduled, policy-managed backup, archive, and space-management facility for file servers and workstations. The term *workstation* will be used in this book to denote workstations, file servers, and PCs.

This publication helps you to get up and running with a usable, practical ADSM system. It explains how to:

- Install the server on Windows NT
- Define the communication network for server and clients
- Customize ADSM for your environment

What You Should Know before Reading This Publication

Before using this publication, you should be familiar with:

- Windows NT operating system
- Devices that will be available to ADSM
- Workstation operating systems on which backup-archive clients reside
- Workstation operating system on which the administrative client resides
- Communication protocols required for the client/server environment

You also need to understand the storage management practices of your organization, such as how you are currently backing up your workstation files and how you are using disk and tape storage.

This section provides an overview of the tasks covered in this book. It also presents an overview of ADSM concepts that are useful in performing those tasks. The steps for each task are described in detail beginning with Chapter 1, "Installing ADSTAR Distributed Storage Manager" on page 1.

Conventions Used in This Book

This section describes the conventions used in this book to make finding and understanding information easier.

Note: If you are viewing this book online, identifying text appears instead of icons.

Types of Information

The following conventions help you to identify types of information.

Important tips or cautions

Detailed information that is useful, but not essential, in performing a task

Tips to help you avoid or correct problems

Where Tasks Are Performed

The following conventions identify where to perform the associated tasks:

Windows NT command line

ADSM Server Utilities

ADSM administrative client command line

ADSM administrative client graphical user interface

ADSM end user client command line

ADSM end user client graphical user interface

How to Use This Book

This book takes you through the tasks for setting up a basic, practical ADSM system. The approach is to take advantage of most defaults supplied by the ADSM commands or graphical user interface.

Installation and Setup Results

Chapter 1, "Installing ADSTAR Distributed Storage Manager" on page 1 describes how to install ADSM. You can choose one of the following setups when you install ADSM: full, compact, or custom. Table 1 shows the difference between **full**, which is recommended for most customers, and **compact**. **Custom** lets you choose the components to install.

Table 1. Full and Compact Installation Types

Full	Compact	Component
X	X	ADSM server installed
X	X	Server options file (dsmserv.opt) created
X		Client options file (dsm.opt) created
X	X	ADSM database created and 5MB volume formatted and defined
X	X	ADSM recovery log created and 9MB volume formatted and defined
X	X	Three storage pools created: BACKUPPOOL, ARCHIVEPOOL, and SPACEMGPOOL
X	X	4MB volume formatted and defined for BACKUPPOOL
X	X	Default storage policy (see "Default Storage Policy" on page viii)
X	X	ADSM device drivers installed
X		Administrative client installed
X	X	Administrative client registered
X		Backup-archive client installed
X	X	Backup-archive client registered
X	X	Two defined backup-archive client schedules (one daily and one weekly incremental backup)
X	X	ADSM server utilities installed
X		ADSM online documentation installed

Chapter 2, "Getting Started After Installation" on page 9 describes how to do the following:

- Start the ADSM Server Utilities and the ADSM server
- Verify the ADSM installation
- Add volumes to your database, recovery log, and disk storage pools
- Add manual tape devices and automatic tape libraries
- Add storage pools for tape volumes
- Arrange disk and tape storage pools in a hierarchy so that client data migrates from disk to tape based on criteria that you specify

Chapter 3, "Customizing Your ADSM System" on page 23 describes how to do following:

- Register new ADSM licenses

- Register additional administrative and backup-archive clients

- Create a new storage policy to meet the needs of a specific group of users in your installation

- Create an include/exclude list template that clients can use to exclude specified data from policy operations

- Define schedules to automate client operations such as backup and archive.

- Define schedules to automate commands issued from the administrative client.

Default Storage Policy

Objects for a default storage policy are defined (all are named STANDARD). This policy specifies the following:

- When a backed up file is no longer associated with a backup copy group, it remains in server storage for 30 days.
- When an archived file is no longer associated with an archive copy group, it remains in server storage for 365 days.
- Client files are not space managed.
- For backup operations:

  - Files are backed up to the default disk storage pool, BACKUPPOOL.

  - An incremental backup is performed only if the file has changed since the last backup.

  - Files cannot be backed up while they are being modified.

  - Up to two backup versions of a file on the client's system are retained in server storage. The most recent backup version is retained for as long as the original file is on the client file system. All other versions are retained for up to 30 days.

  - One backup version of a file that has been deleted from the client's system is retained in server storage for 60 days.

  - For archive operations:

- Files are backed up to the default disk storage pool, ARCHIVEPOOL.

- Files cannot be archived while they are being modified.

- An archive copy is kept for up to 365 days.

Basic Tasks

This section presents an overview of the tasks of installing, setting up, and customizing ADSM. The following list names each task and points to the details of the procedures:

1 Installing ADSM, which includes the server, one Windows 32-bit administrative client and one Windows 32-bit backup-archive client (page 1). See "Setting Up the Windows 32-bit Client" on page 46 for communications information for the Windows 32-bit client.

Information on installing additional clients can be found in *ADSM Installing the Clients*.

2 Getting started after installation (page 9)

   a Starting the ADSM Server Utilities (page 9)

   b Starting the server and device driver (page 9)

   c Verifying the ADSM installation (page 11)

   d Increasing the size of the database and recovery log (page 11)

   e Increasing the size of a disk storage pool (page 11)

   f Starting and ending an administrative client session (page 12)

   g Using tape devices with ADSM (page 13)

   h Stopping the server (page 22)

3 Customizing ADSM (page 23)

   a Registering ADSM licenses (page 23)

   b Registering additional administrators (page 24)

   c Registering additional client nodes (page 25)

   d Creating new policies (page 25)

   e Creating schedules to automate client operations (page 30) and administrative commands (page 33)

Advanced Tasks

To take greater advantage of the power of ADSM, you may want to perform other tasks described in the *ADSM Administrator's Guide*, including the following:

Server and storage pool protection
You can:

- Mirror your database and recovery log
- Back up your database and storage pools
- Recover lost or damaged database and storage pool data

Drives and libraries
You can:

- Add automated libraries to ADSM
- Support ADSM with your media management system

Storage Pools and Volumes
You can:

- Enhance your storage pool definitions
- Move files from one volume to another

Storage Policies
You can enhance your storage policy definitions to:

- Customize policies for groups of clients
- Provide clients on some platforms with space management services
- Provide additional functions to backup and archive services

Automation of Client Operations and Administrative commands
You can:

- Enhance your schedule definitions
- Optimize the timing of scheduled operations

Security
You can:

- Let clients access ADSM without requiring passwords
- Modify the password expiration interval
- Assign levels of administrator authority

Server Operations
You can monitor and control:

- Server processes
- Client sessions
- ADSM activity log
- ADSM accounting records

Server Data
You can copy all or part of the server data to removable media and transfer it to a second server.

Before You Continue

Here are some questions that you should answer before you start to install and customize your ADSM system:

- Which workstations in my network will require ADSM services and what communications method will they use to attach to the server?

- Who will I need to administer ADSM?

- What devices do I want to include in my ADSM system?

    – What are the device types of each?

    – What are the capacities of the device media?

    – Which devices are in automated libraries and which require operator intervention?

- What are the backup, archive, and space management needs of the ADSM users? Are there groups that have different needs? Within groups should some files be managed differently? For example:

    – How often should client files be backed up?

    – How many versions of client files should be retained?

    – How long should those versions be retained?

    – What types of data require quick access and what types do not?

- Which client operations do I want to automate and when should they be scheduled?

- Which administrative commands do I want to automate and when should they be scheduled?

ADSM Concepts

The rest of this section describes some basic ADSM concepts that can help you in your installation and customization. If you are already familiar with ADSM, you can go directly to Chapter 1, "Installing ADSTAR Distributed Storage Manager" on page 1.

How ADSM Can Change the Way You Work

This book introduces you to a different way of managing your data storage. ADSTAR Distributed Storage Manager (ADSM) provides centrally administered services based on a set of defined policies. Many services can be automated and scheduled to make the best use of your data processing resources.

After you tailor your ADSM system to the needs of your installation, it will run with a minimum of intervention from you. Your initial effort will be to define to ADSM the devices and media it has access to, the clients it will supply services to, the policies that will govern those services, and schedules for automating your operations.

The benefits of this effort include increasing your productivity, eliminating error-prone, repetitive tasks, and providing a flexible, consistent set of policies to govern many types of data owned by a wide variety of user groups.

This section introduces you to some ADSM concepts that will help you perform the tasks described in the rest of the book. It also includes an overview of those tasks and the more advanced tasks described in the ADSM Administrator's Guide. Finally, this section provides some questions to consider when you are planning the installation and customization of ADSM.

System-Managed Storage

ADSTAR Distributed Storage Manager (ADSM) includes the following components:

Server
Provides ADSM services to client workstations. The server maintains a database and recovery log for ADSM resources, users, and user data.

Server Storage
Consists of storage pools of random and sequential access media. The storage pools contain files that are backed up, archived, and migrated from client nodes.

Server Utilities
Provides an interface to help configure the ADSM server.

Administrative client
Provides an interface for controlling the ADSM server.

Backup-archive client
Provides backup-restore and archive-retrieve services to workstations and file servers.

Hierarchical storage management (HSM) client
Provides space management services for workstations on some platforms.

Server Storage Migration

One goal of system-managed storage is to ensure the most efficient use of your storage resources. You can structure server storage into a hierarchy. You can then define criteria by which data is migrated from faster, more expensive storage pools (disk, for example) to slower, less expensive storage pools (tape, for example).

ADSM Policy

ADSM policy governs storage management, including:

Backup
Copying files from client workstations to server storage to ensure against loss of data. Copies of multiple versions of a file can be stored.

Archiving
Copying files from client workstations to server storage for long-term storage.

Space Management

Freeing client storage space by copying a file from client workstations to server storage. The original file is replaced with a much smaller file that points to the location of the original in server storage. This is also called **migration**.

Policy is defined by administrators in policy objects: policy domains, policy sets, management classes, and backup and archive copy groups.

Figure 1 shows an overview of the ADSM process for storing client data.
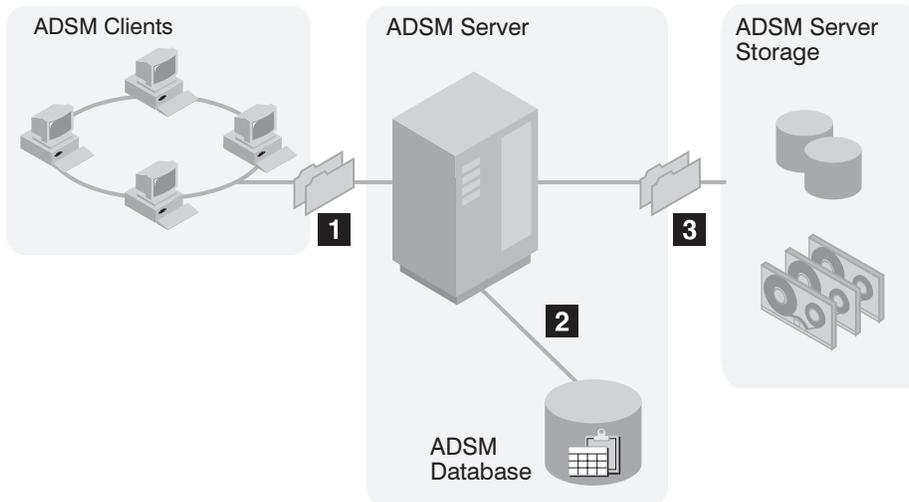
Figure 1. Overview of How ADSM Stores Client Data

When clients back up, archive, or migrate files, ADSM:

1 **Determines where to store the file**

ADSM checks the management class bound to the file to determine the destination of the file, that is, where the file should be stored. A destination might be a group of disk or tape volumes. These groups of volumes are called storage pools. Copy groups, which are within management classes, specify destinations for backed up and archived files. Management classes specify destinations for space-managed files.

2 **Stores information about the file in the ADSM database**

ADSM saves information in the database about each file in server storage. For example, ADSM saves the name and location of the file.

3 **Stores the file in ADSM server storage**

ADSM stores the client files in disk or tape storage pools.

Figure 2 on page xiv shows the interaction between ADSM policy objects and ADSM backup, archive, and space management services.
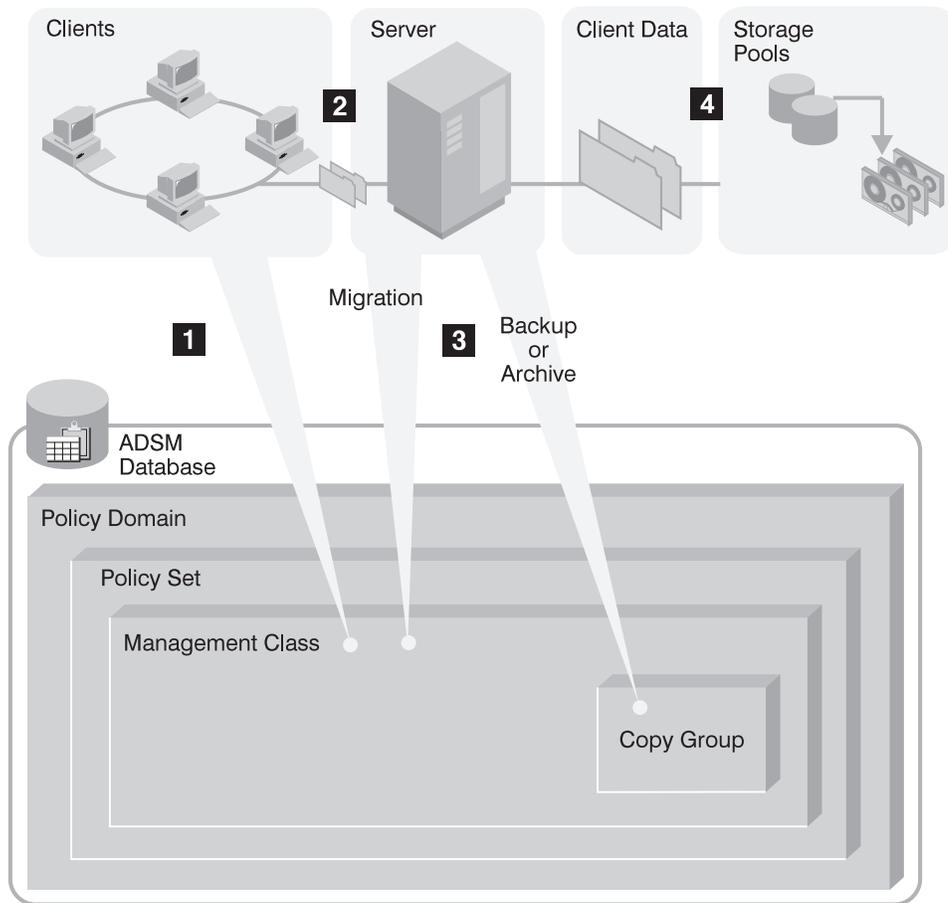
Figure 2. How ADSM Controls Backup, Archive, and Space Management

1 An ADSM client backs up, archives, or migrates a file. The file is bound to either the default management class or a management class specified in the client's include-exclude list.

2 If, according to the management class, the file is eligible for backup, archive, or space management, the client sends the file and file information to the server.

3 The server checks the management class or copy group to determine where in server storage to store the file initially.

4 The server stores the file in the appropriate storage pool and information about the file in the database.

When files in server storage are migrated from one pool to another, the server updates the file information in the database.

Expiration

Files remain in server storage until they expire and expiration processing occurs. A file expires based on criteria set in ADSM policy. For example, backup policy may specify that up to two backup versions of a file can be kept in server storage. If a third backup version of a file is created, the oldest version in server storage expires. Backup policy may also specify that if a file is deleted from a client file system, it expires after 60 days.

During expiration processing, ADSM deletes from the database those files that have expired.

Storage Devices and Media

ADSM represents physical storage devices and media with the following administrator-defined storage objects:

- Device class
- Library
- Drive
- Storage pool
- Storage volume

Device Class

Each device is associated with a device class that specifies the device type and how the device manages its media. ADSM has a predefined device class (DISK) for random access devices. Administrators can also define device classes for sequential access devices, such as 8mm tape drives. If the sequential device is a tape drive, the device class is associated with a library and a drive.

Library

An ADSM library is an administrator-defined collection of one or more drives (and possibly robotic devices) that share similar media mounting requirements.

Drive

Each drive mechanism within a tape device is represented by an ADSM drive. For devices with multiple drives, including automated libraries, each drive is separately defined to ADSM.

Storage Pools

A storage pool is a named collection of storage volumes that are associated with one device class. Each storage pool represents a collection of volumes sharing the same media type. For example, an 8mm tape storage pool contains only 8mm tape volumes.

Storage Pool Volumes

A storage pool volume is associated with a specific storage pool and represents a unit of space available for ADSM client data.

Figure 3 on page xvii summarizes the relationships among the physical device environment, ADSM storage objects, and ADSM clients. The numbers below correspond to the numbers in the figure.

1   When clients are registered, they are associated with a policy domain. Within the policy domain are the other ADSM policy objects.

2 , **3**

When a file is backed up, archived, or migrated from a client, it is bound to a management class. A management class and the backup and archive copy groups within it specify where files are stored and how they are managed when they are backed up, archived, or migrated from a client (space-managed files).

4 , **5**

Storage pools are the destinations for backed up, archived, or space-managed files. Copy groups specify storage pools for backed up or archived files. Management classes specify storage pools for space-managed files.

Storage pools are mapped to device classes, which represent devices. The storage pool contains volumes as indicated in the device type associated with the device class. For example, a storage pool that is mapped to a device class with a device type of 8MM contains only 8mm tapes.

All devices require a device class that specifies a device type. Tape devices also require a library and drive for management of media, including the mounting of that media.

6   Files that are initially stored on disk storage pools can migrate to tape storage pools if the pools are set up in a storage hierarchy.

Clients

1

Policy Domain

4

Disk

Policy Set

Volume   Volume

· · ·

Management Class

2   Copy
Group   **Points to**

Storage
Pool   **Represents**

DISK
Device Class

Migrate   **6**

Tape

5

Volume   Volume

· · ·

Management Class

3   Copy
Group   **Points to**

Storage
Pool   **Represents**   Media

Device Class

Library   **Represents**

Drive   Drive

· · ·

Drives

Device
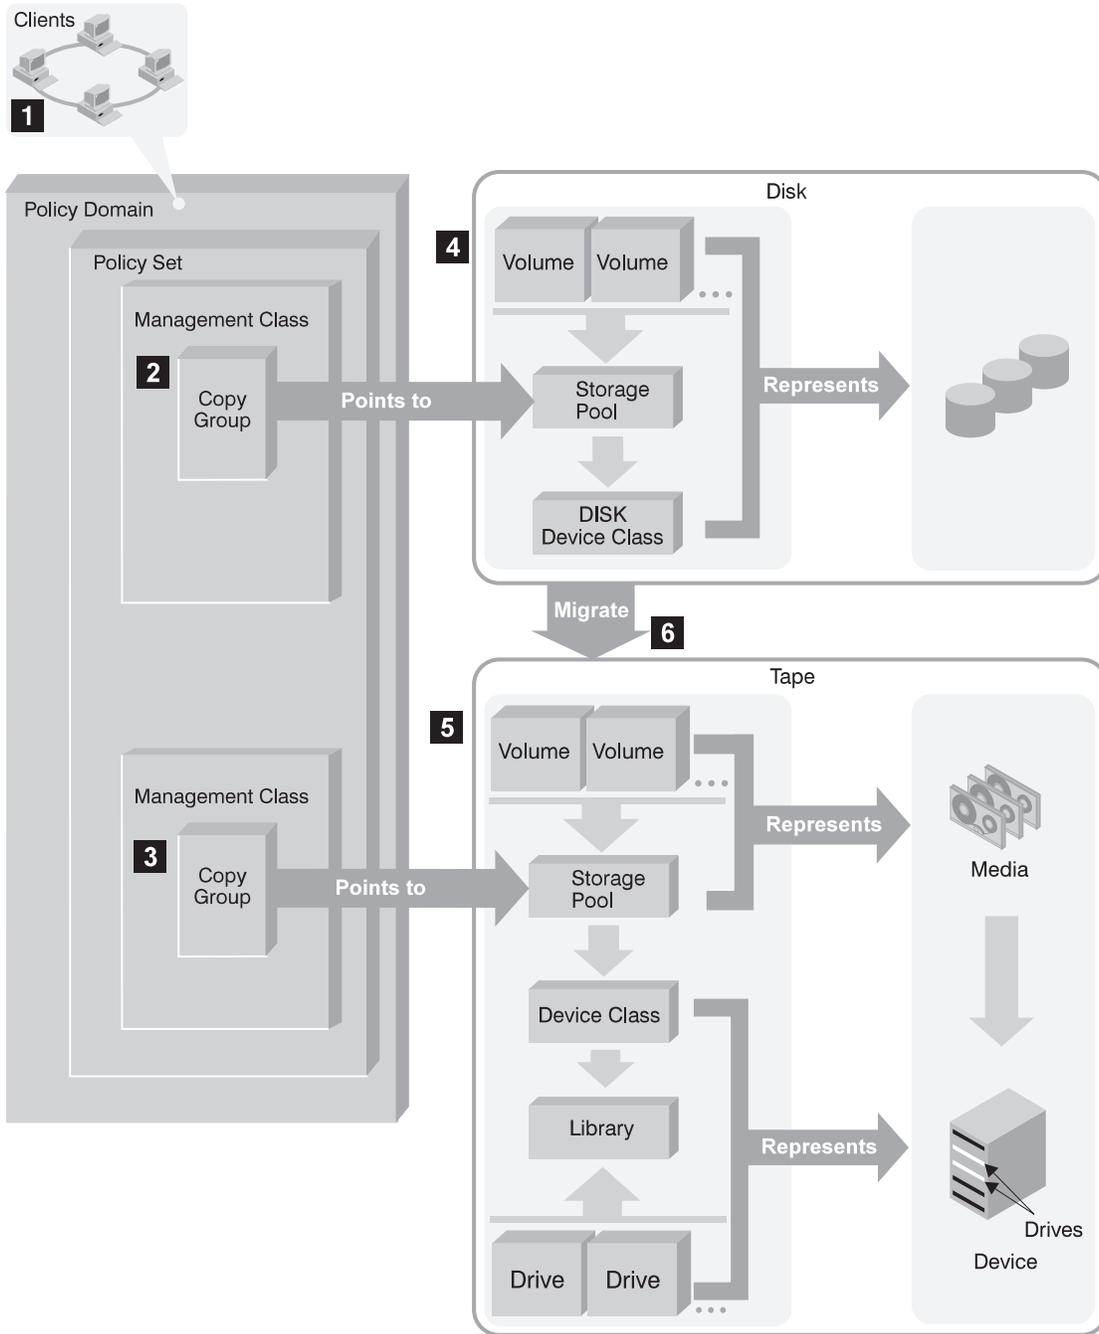
Figure 3. Putting it All Together

Automating Client Operations

You can automate client operations, such as backup and archive. Figure 4 on page xix shows the ADSM objects that may be involved in automated client operations. They are:

Include-exclude criteria on each ADSM client
Determines which files are backed up, archived, or space-managed, and binds files to management classes

Management class
Determines where client files are stored and how they are managed

Schedule
Determines when client operations such as backup occur

Association defined between client and schedule
Determines which schedules are run for a client
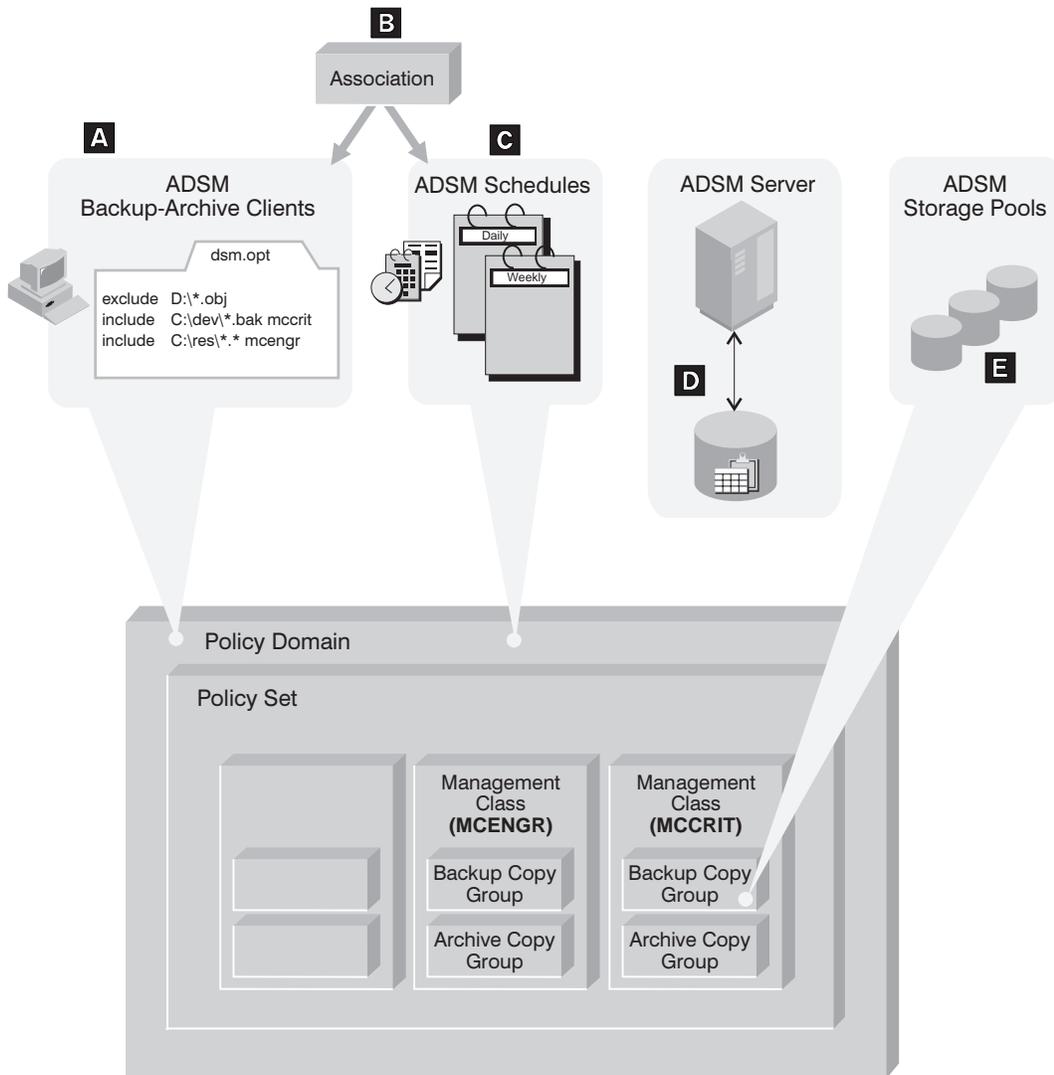
Figure 4. Automating Client Operations

The client can specify a management class for a file or set of files, or can use the default management class for the policy domain. The client specifies a management class by using an INCLUDE option in the client's include-exclude list or file. (See **A** in Figure 4.)

The management class contains information that determines how ADSM handles files that clients backup, archive, or migrate. For example, the management class contains the backup copy group and the archive copy group. Each copy group points to a *destination*, a storage pool where files are stored when they are backed up or archived. (See **E** in Figure 4.)

Clients are assigned to a policy domain when they are registered. Schedules that can automate client operations are also associated with a policy domain. (See **C** in Figure 4.) To automate client operations, you define schedules for a domain. Then you define associations between schedules and clients in the same domain. (See **B** in Figure 4 on page xix.)

For a schedule to work on a particular client, the client machine must be turned on and must be running the client scheduler.

The scheduled client operations are called *events*, and information about events are stored in the ADSM database. (See **D** in Figure 4 on page xix.) For example, you can query the server to determine which scheduled events completed successfully and which failed.

Chapter 1. Installing ADSTAR Distributed Storage Manager

The ADSM installation program includes the server, the server utilities, the administrative client, the 32-bit Windows client, and the online books.

During installation, a default administrative client and backup-archive client are registered with the server as follows:

- Administrative client
  ID name   admin
  Password   admin
- Backup-archive client
  ID name   client
  Password   client

The default ADSM directory is **\win32app\ibm\adsm**.

!

Before You Install

Ensure that:

- You have a Microsoft Windows NT 3.51 or later workstation with at least 100MB of free disk storage (for a full installation) and 16MB (32MB is recommended) of memory.

- You have a CD-ROM device available for installation.

- You are logged on to Windows NT with a user name having administrator authority.

Installing the ADSM Server, Clients, and Utilities

1 Insert the ADSM for Windows NT CD #3 into your CD-ROM drive.

2 Perform one of the following tasks:

WINDOWS NT FILE MANAGER

1. Start **File Manager**
2. Double-click on the file **setup.bat** located in the root directory of your ADSM CD #3

From the root directory of your CD #3, type:

```
setup
```

3 Read the **Welcome** screen and press **Next** to continue.

4 Click on one of the **Setup Options** from this screen:
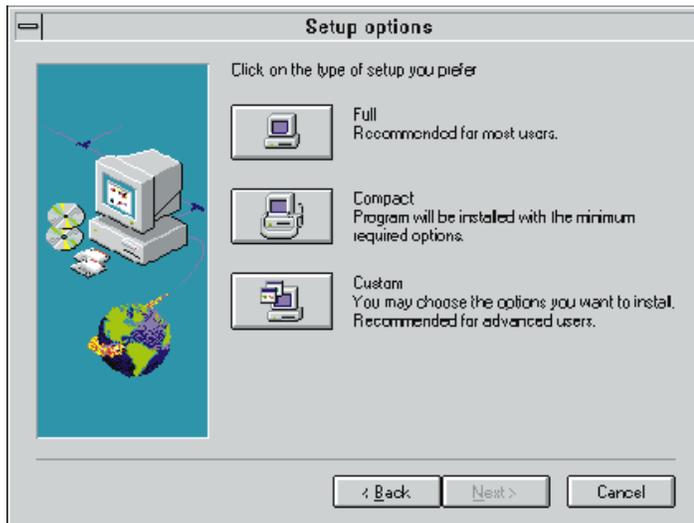
Figure 5. Setup Options Screen

- FULL (Recommended)

  See "Installation and Setup Results" on page vii for the results of this installation. The FULL setup requires approximately 100MB disk space.

- COMPACT

  This selection installs only the server component. The ADSM Utilities are included in the server component. This selection requires approximately 60MB disk space.

- CUSTOM

  This selection allows you to choose which components to install.

  Note: This example covers a FULL setup.

5 A **Select Destination Directory** screen appears. Keep the directory that is shown in the window or change the path in which to install ADSM.
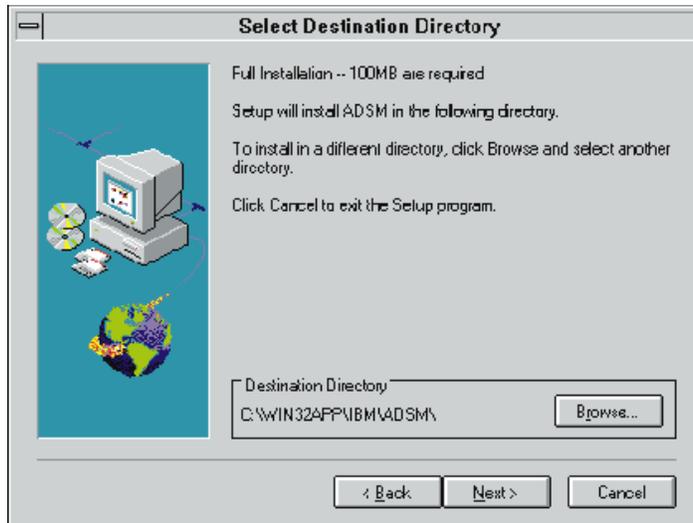
Figure 6. Select Destination Directory Screen

> **Note:** You can change the path in which to install ADSM by clicking on the **Browse** button at the bottom right of this screen. The destination directory that you choose becomes the ADSM root directory. Subdirectories for each component are created from this root. For example, if you install to **c:\win32app\ibm\adsm**, the server is placed in **c:\win32app\ibm\adsm\server**.

6 Click on **Next**.

A **Select Program Folder** window appears. Type in or select the name of a folder to be used to hold the ADSM components.
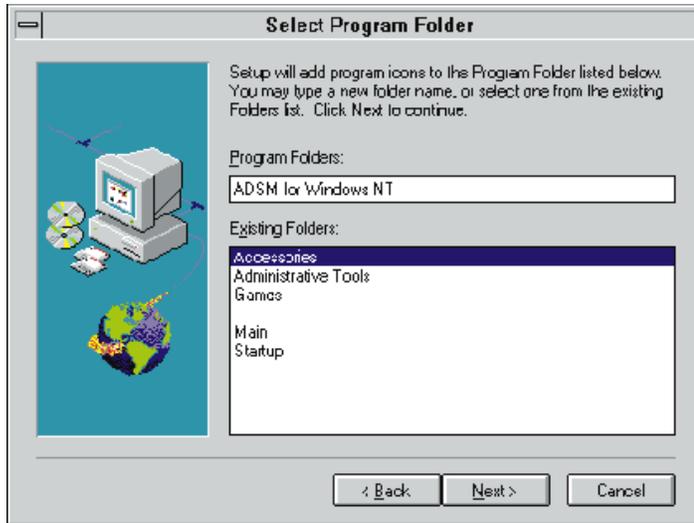
Figure 7. Select Program Folder Screen

7 Click on **Next**.

The **Start Copying Files** screen appears. Review the current settings.
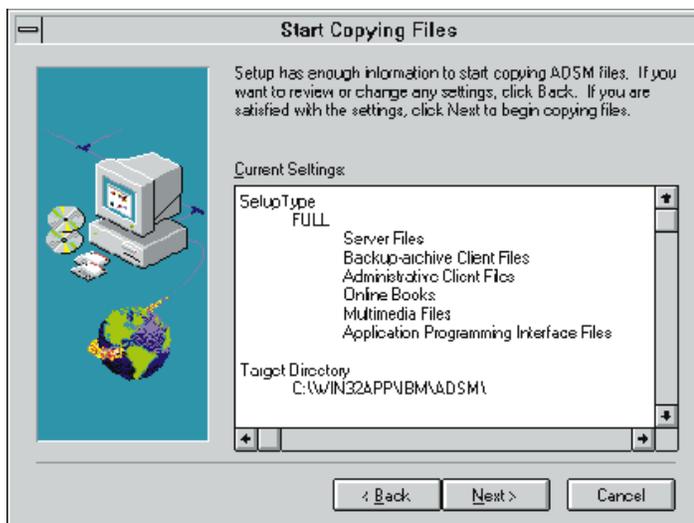
Figure 8. Start Copying Files Screen

8 Click on **Next**.

A progress screen appears showing the files being installed. While the files are being installed, these events occur:

- A program group is created (see Figure 10 on page 7).

- The ADSM server database and recovery log are initialized. An NT command window provides information on what is occurring. The results of this step are logged to the file **initserv.log** in the server directory.

- The ADSM server initial configuration is applied. Again, an NT command window appears while the configuration completes. The results of the configuration are logged to **cfgserv.log** in the server directory.

  As part of the initial configuration:

  – The default administrative client is registered with:
     ID name   admin
     Password   admin
     Authority   system
  – The default backup-archive client is registered with:
     ID name   client
     Password   client

- A 4MB disk storage volume is created and defined to the ADSM BACKUPPOOL disk storage pool.

- A daily and weekly incremental backup schedule is defined. After the nodes are registered with the server, you can associate them with schedules. You can define new schedules or modify these default schedules.

! If you encounter problems starting the ADSM server after installation, go to your server directory and examine these two files:

- initserv.log
- cfgserv.log

You may be able to fix the problem and continue. If you do require technical support, you may be asked for these files for problem determination.

---

9 When the process completes, a **Setup Complete** screen appears.

Figure 9. Setup Complete

You have these choices:

- Yes, I want to restart my computer now

  If you select Yes, the computer will be rebooted. See "Start the Server" on page 9 to start the server after reboot.

  If you select Yes, you will need to click on the the **Getting Started** icon to review the **ADSM Readme** files.
- No, I will restart my computer later

  If you select No, you see another panel with the administrative and client names and passwords.

When you reboot the computer, Windows NT recognizes the ADSM device driver and includes ADSM in event logging.

10 Click on **Finish** with this line checked or not:

```
Yes, I want to run the server utilities now
```

This screen also has a note saying that all ADSM components can be launched from the ADSM Server Utilities.

The **Getting Started** help file appears (unless you rebooted in Step 9 on page 5) so that you can review the ADSM Readme files and see a product overview.

Note: The ADSM program screen shown here is for Windows NT 3.51. If you are using Windows NT 4.0, this will be arranged as a cascading menu of ADSM icons and labels.

Figure 10. ADSM for Windows NT Screen

11 Your installation is complete.

See the *ADSM Administrator's Reference* for instructions on installing additional ADSM servers on a Windows NT system with the DSMSERV INSTALL command.

Chapter 2. Getting Started After Installation

This section describes how to start your ADSM server, verify your installation, and prepare ADSM for production.

> **!** This section includes the procedure for setting up and using tape devices with ADSM. One step is to start the device driver. To do this step, you must first reboot your machine after installing ADSM.

Further customization is described in Chapter 3, "Customizing Your ADSM System" on page 23.

Before you begin, you should be aware of some online sources of information:

Getting Started   Select **Getting Started** from the utilities help menu. This gives you access to the **ADSM Readme** files and provides a graphical introduction to ADSM.

Online books   Click on the **Online Books** button on the Utilities toolbar to access the ADSM online books.

ADSM Home Page   Click on the **Web** button on the Utilities toolbar to start a web browser and go to ADSM-related pages. You must have Netscape, Mosaic, or MS Internet Explorer to use this utility.

Step 1: Start the ADSM Server Utilities

Double-click on the **ADSM Server Utilities** icon:

Step 2: Start the Server and the Device Driver

The ADSM server must be running before any client can connect, and the device driver must be running before you can use a device with ADSM.

Start the Server

There are several ways to start the server:

- If you start the server as a service from the ADSM Utilities, the server remains active when you log off the workstation.

  From the ADSM Utilities, do the following:

1. Double-click on the **Services** icon.

2. Click on the **ADSM Server** tab.

3. If the status is **Stopped**, click on the **Start** push button.

Note: If you have a problem starting the server, you can click on the **QA** button. Common problems and their solutions are described.

- From the ADSM for Windows NT program folder, double-click on the **ADSM Server** icon.

Start the Device Driver

Ensure that your devices are physically attached to your system according to the manufacturer's instructions. Next start the device driver. The ADSM device driver, ADSMSCSI, is installed on your system as part of the server installation. The ADSMSCSI driver must be used for all tape autochangers. If you want to use another driver for a manually operated tape device, see the *ADSM Administrator's Guide* for directions. The examples in this chapter assume the use of the ADSMSCSI driver.

Start the device driver from the ADSM Utilities:

1. Double-click on the **Services** icon.

2. Click on the **ADSM Device Driver** tab.

3. Click on the **Start** push button.

Note: If you have a problem starting the device driver, you can click on the **QA** button. Common problems and their solutions are described.

Step 3: Verify the ADSM Installation

To verify your installation, do the following:

1 To start the backup-archive client GUI, double-click on the client icon:

The node name and password are set to defaults in the client options file (dsm.opt).

2 Back up several files.

A drive must be labeled before files on that drive can be backed up.

For example, to back up all the files in the root directory of a disk drive (for example, c:\), do the following:

1. From the **Drives for Backup/Archive** section of the **Drive Information**, select drive c: by clicking on the drive
2. Click on **ADSM**
3. Click on **Backup by File specification** from the pull-down menu
4. From the **Backup by File Specification** window, enter ***.***
5. Click on the **List Files** push button to see the selected files
6. Click on the files you want from this list, or click on the **Select All** push button
7. Click on **Backup** push button

END USER

! ADSM will issue a message indicating whether or not the backup was successful. If it was unsuccessful, you can go to the README file for help with common problems (select **Getting Started** from the utilities help menu).

Step 4: Increase Database and Recovery Log Size

At installation, a 5MB database volume (db1.dsm) and a 9MB recovery log volume (log.dsm) are created. The database size is largely determined by the number of client files being stored on server storage. As you add clients, you may need to increase your database size. The recovery log size may need to be increased depending on the current utilization.

To create additional database and recovery log volumes for the server, do the following:

1. From the ADSM Utilities window, double-click on the **Formatter** icon.
2. From the Formatter window, click on the database volumes push button or the recovery log volumes push button. This brings up a wizard that estimates the size of additional volumes and creates those volumes.

Step 5: Increase Disk Storage Pool Size

At server installation, three default disk storage pools are created: BACKUPPOOL, ARCHIVEPOOL, and SPACEMGPOOL. BACKUPPOOL has one 4MB volume (DATA1.DSM) defined to it but should be at least large enough to hold one day's worth of client backups. No other volumes are defined. You can use a wizard in the ADSM Utilities to create and define additional volumes for any storage pool at any time.

To increase the size of the BACKUPPOOL and to create and define a volume for ARCHIVEPOOL, do the following:

1. From the ADSM Utilities window, double-click on the **Formatter** icon.

2. From the Formatter window, click on the disk volumes push button. This brings up a wizard that estimates the size of additional volumes and creates those volumes.

Step 6: Start and End an Administrative Client Session

This section describes how to start and end a Windows 32-bit administrative client session. For information on starting all other clients, see the *ADSM Installing the Clients*.

When you start the administrative client, you must enter an administrator name and password. The defaults for the administrative client created at ADSM installation are:

Name   *admin*
Password *admin*

To change these values see "Change Administrative Client Passwords" on page 24.

1 Start the Windows 32-bit **administrative client GUI** by double-clicking on the administrative GUI icon:
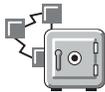
2 Start the Windows 32-bit **administrative client command line interface**, by double-clicking on the administrative command line client icon:

ADMIN

To end an administrative client session do the following:

From a GUI client, one of the following:

• Double-click on the system menu symbol

OR

ADMIN

• Click on the system menu symbol, and click on **Close** from the pull down menu

From a command line client in interactive mode:

`quit`

ADMIN

Step 7: Use Tape Devices with ADSM

Before you can use tape with ADSM, you must do the following:

1. Start the device driver

2. Determine the device names

3. Define a library for the drives, the drives, and a device class and its associated storage pool

4. Include the storage pool in your hierarchy

To perform these tasks, use the ADSM Utilities, administrative command line, and administrative GUI.

The following terms may help you to understand the task you are to perform:

device class A group of storage devices of the same type, for example, disk or 8mm.

library A collection of one or more drives with similar mount mechanisms. For ADSM to use it, a device must be assigned to a library. There are two types of libraries:

  manual Volumes are loaded by a human operator.
  automated Volumes are loaded by a robotic autochanger device.

This section describes the following tasks on the pages indicated:

Task	Page
Determine the device names	15
Add a manual tape drive	15
Add an automated tape library	16
Prepare tape volumes for use	18
Verify the use of the tape devices.	21

If you want to use a tape management system with ADSM, see the *ADSM Administrator's Guide* for details.

The examples in this section set up the following configuration:

- A manual library consisting of two Exabyte 3500 8mm tape drives

- An Exabyte EXB-210 automated tape library with two Exabyte 3500 8mm tape drives

- A tape storage pool associated with these libraries. All the ADSM disk storage pools migrate data to this tape storage pool.

Here are some concepts that can help you to set up your tape environment:

Storage pool migration

To prevent disk storage pools from becoming filled with client data, you can specify that client data be automatically migrated to the tape storage pool.

A disk's *high migration threshold* (default value: 90%) determines when data is migrated. The *low migration threshold* (default value: 70%) Determines when migration stops. Thus data migration from the default backup disk storage pool begins when the pool is 90% full and continues until it is at 70% full or less.

Another default is to cache data (leave the data on the storage pool) even after the data has been migrated. Cached data stays on disk for quick access until the space is needed for new data.

See *ADSM Administrator's Guide* for details.

Mount limit

Mount limit (the number of drives available in a device class) has a default value of 1. The mount limit should be equal to the number of drives of the same type in that library.

Scratch volumes

A scratch volume is a volume that is available for ADSM use. The volume is labeled, is either blank or contains no valid data, and is not defined to ADSM. You can specify the number of scratch volumes that ADSM can request for this pool (MAXSCRATCH parameter in the command or the **Maximum scratch volumes** field in the GUI). If you do not specify a value, no scratch volumes can be requested, and you must define each volume to be used.

Collocation

This is process by which the server attempts to keep all files belonging to a client on a minimal number of tape volumes. Collocation is turned off by default. Once clients begin storing data in a storage pool with collocation off, you cannot easily change the data in the storage pool so that it is collocated. To understand the advantages and disadvantages of collocation, see the *ADSM Administrator's Guide*.

Determine Device Names

Determine the device names of your tape drives and autochangers. Devices controlled by the ADSMSCSI driver have names with the form mt*x.y.z.n* or lb*x.y.z.n*, where:

mt   Tape device
lb   Indicates an automated library
x   Device SCSI ID
y   Device logical unit number (LUN)
z   Port number for the SCSI adapter device driver
n   Bus number supported by the adapter device driver

For example, mt3.0.0.0 represents the tape drive at SCSI ID 3, LUN 0 on port 0, bus 0.

Ensure that the device driver is started (see "Start the Device Driver" on page 10 for details). The Diagnostics utility searches for devices and displays information about them:

From the ADSM Utilities window, double-click on the **Diagnostics** icon and then click on the **Refresh** button. The Device Detector displays the device name and ADSM device type of each supported device.

Add a Manual Tape Library

The commands for defining a manual tape library are, in the order issued:

1. DEFINE LIBRARY
2. DEFINE DRIVE
3. DEFINE DEVCLASS
4. DEFINE STGPOOL

1 Define a manual library and two drives associated with it.

The library is named MANUALLIB. The two drives, named TAPE01 and TAPE02, are at SCSI IDs 1 and 2.

ADMIN

```
define library manuallib libtype=manual
define drive manuallib tape01 device=mt1.0.0.0
define drive manuallib tape02 device=mt2.0.0.0
```

2 Classify the drives by device type. It is helpful to choose a device class name that identifies the type of device for which it is intended (in this example, TAPECLASS).
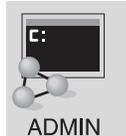
```
define devclass tapeclass devtype=8mm
  library=manuallib mountlimit=2
```

ADMIN

3 Define a storage pool named TAPEPOOL for the device class.

```
define stgpool tapepool tapeclass maxscratch=100
```

ADMIN

OR

1. Double-click on **Storage Pools**
2. Click on **Edit** from the menu bar
3. Click on **Add Primary Storage Pool** from the pull-down menu
4. In the **Storage pool name** field, enter TAPEPOOL
5. Click on the **Device class** tab, and enter the following:

   Device class TAPECLASS (from the drop-down list)
   Maximum scratch volumes 100

6. Click on the **Add** push button

ADMIN

You can also use the DEFINE STGPOOL command to perform the preceding task.

Add an Automated Tape Library

The commands for defining an automated tape library are, in the order issued:

1. DEFINE LIBRARY
2. DEFINE DRIVE
3. DEFINE DEVCLASS
4. DEFINE STGPOOL

1 Define an automated tape library and its two internal drives.

The library is an Exabyte EXB-210, named AUTOLIB, whose SCSI library mechanism is at SCSI ID 3. The two drives, named TAPE05 and TAPE06, are at SCSI IDs 4 and 5. The element numbers for the drives in an EXB-210 are 82 and 83.
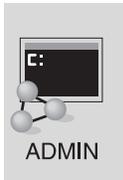
```
define library autolib libtype=scsi device=lb3.0.0.0
define drive autolib tape05 device=mt5.0.0.0
  element=82
define drive autolib tape06 device=mt6.0.0.0
  element=83
```

ADMIN

2 Classify the drives by device type. It is helpful to choose a device class name that identifies the type of device for which it is intended (in this example, AUTOCLASS).
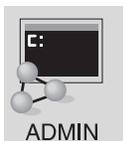
```
define devclass autoclass devtype=8mm
  format=8500 library=autolib mountlimit=2
```

ADMIN

3 Define a storage pool named AUTOPOOL for the device class.

```
define stgpool autopool autoclass
  maxscratch=100
```

ADMIN

OR

1. Double-click on **Storage Pools**

2. Click on **Edit** from the menu bar

3. Click on **Add Primary Storage Pool** from the pull-down menu

4. In the **Storage pool name** field, enter AUTOPOOL

5. Click on the **Device class** tab, and enter the following:

   Device class   AUTOCLASS (from the drop-down list)
   Maximum scratch volumes 100

6. Click on the **Add** push button

ADMIN

You can also use the DEFINE STGPOOL command to perform the preceding task.

Display Device Information

Before continuing, you can view the results of your definitions. If necessary, you can update these definitions by using update commands.

- To query the libraries:

  `query library`
- To query the drives:

  `query drive`
- To query the device classes:

  `query devclass format=detailed`
- To query the storage pools:

  `query stgpool`

ADMIN

Prepare Tape Volumes

All tape volumes must have IBM standard labels before ADSM can use them. You can use tapes as scratch volumes, up to the number of scratch volumes specified for the storage pool. Using scratch volumes allows ADSM to acquire volumes as needed.

You can also use private volumes in a storage pool. You must define any private volumes to ADSM, assigning each to a specific storage pool. However, if a storage pool contains only private volumes and runs out of them, storage operations for that pool stop until more volumes are defined.

Prepare Tape Volumes with a Manual Drive

If you are using only one tape drive, consider labeling several tapes at this time.

> ! You cannot label tapes while the server is writing to or reading from the drive.

1 If the tape volumes do not already have IBM standard labels, label them.

  1. Insert the new volume into drive TAPE01 (device mt1.0.0.0).

  2. From the ADSM Utilities main screen, select **Labeler**

  3. Click on the manual devices button

  4. Enter a 1 to 6 character volume name (for example, DSM001).

  5. Select the device name from the list of devices

  6. Click on the **Apply** push button

(for example, DSM001).

2 Define any volumes that are to be used as private volumes for example, DSM001. For example, define the volume you just named DSM001.

1. Expand **Storage Pools**

2. Double-click on **Storage Pool Volumes**

3. Click on **Edit** from the menu bar

4. Click on **Add** from the pull-down menu

5. Enter the following:

ADMIN

Storage pool name	TAPEPOOL (from the drop-down list)
Volume name	DSM001

6. Click on the **Add** push button

You can also use the DEFINE VOLUME command to perform the preceding task.

Prepare Tape Volumes with an Automated Library

To label tapes with an automated library, remove any tapes that you do not want to use with ADSM and load the tapes to be labeled.

1 If the tape volumes do not already have IBM standard labels, label them.

1. From the ADSM Utilities main screen, select **Labeler**

2. Click on the autochanger devices icon

3. Type a 1 to 6 character volume label (for example, DSM001).

4. Type the element number for the device

5. Click the **Search** and **Keep** check boxes

6. Select the autochanger and drive from the list of devices

7. Click on the **Finish** push button

8. You are prompted to enter a 1 to 6 character volume label for each additional tape found

2 Check the tape volumes into the automated library.

Ensure that the volumes are in their slots, and enter:

```
checkin libvolume autolib status=scratch search=yes
```

ADMIN

The library checks in each volume automatically.

3 Define any volumes that are to be used as private volumes (for example, DSM001).

1. Expand **Storage Pools**

2. Double-click on **Storage Pool Volumes**

3. Click on **Edit** from the menu bar

4. Click on **Add** from the pull-down menu

5. Enter the following:

Storage pool name	AUTOPOOL (from the drop-down list)
Volume name	DSM001

6. Click on the **Add** push button

You can also use the DEFINE VOLUME command to perform the preceding task.

Include Tape in the Server Storage Hierarchy

Now that ADSM can use your tape drives, you can update your disk storage pools so that client data can be migrated to tape. This section describes how to change BACKUPPOOL so that data migrates to AUTOPOOL and how to change ARCHIVEPOOL so that data migrates to TAPEPOOL

1. Double-click on **Storage Pools**

2. Click on BACKUPPOOL

3. Click on **File** from the menu bar

4. Click on **Open as properties** from the pull-down menu

5. Click on the **Migration(2 of 2)**, and for **Next storage pool**, do the following:

   a. Deselect (by clicking on) the **None** check box
   b. Click on AUTOPOOL from the **Next pool** drop-down list

6. Click on the **Apply** push button

7. Click on ARCHIVEPOOL

8. Click on **File** from the menu bar

9. Click on **Open as properties** from the pull-down menu

10. Click on the **Migration(2 of 2)**, and for **Next storage pool**, do the following:

    a. Deselect (by clicking on) the **None** check box
    b. Click on TAPEPOOL from the **Next pool** drop-down list

11. Click on the **Apply** push button

You can also use the UPDATE STGPOOL command to perform the preceding task.

Verify the Tape Device Environment

Do the following to verify that the tape devices have been integrated into your ADSM system. These steps assume that one or more clients have backed up data to a disk storage pool. Data in the disk storage pool is then migrated to a tape storage pool.

1 Prepare for tape mount requests.

Migration from a disk to a tape storage pool requires tape mounts. The mount messages are directed to the server console and to any administrative client that has been started with either the mount mode or console mode option. You must monitor these messages.

If you are running the server as a service, start an administrative client in mount mode to monitor the mount messages. In the ADSM Server Utilities, choose **Launch**, then choose **Administrative Client as Mount Monitor**.

2 Migrate data from a disk storage pool.

Trigger migration from a disk storage pool (BACKUPPOOL, for example) by setting the high and low migration thresholds to 0. After migration occurs, reset the thresholds to their original settings.

1. Double-click on **Storage Pools**

2. Click on BACKUPPOOL

3. Click on **File** from the menu bar

4. Click on **Open as properties** from the pull-down menu

5. Click on the **Migration(2 of 2)**, and enter the following:

   Migration Thresholds Low   0
   Migration Thresholds High   0

ADMIN

6. Click on the **Apply** push button

You can also use the UPDATE STGPOOL command to perform the preceding task.

ADSM requests tape mounts as needed and moves any data in the disk storage pool to the volumes in the tape storage pool.

3 After migration occurs, reset the thresholds to their original settings.

1. Double-click on **Storage Pools**

2. Click on BACKUPPOOL

3. Click on **File** from the menu bar

4. Click on **Open as properties** from the pull-down menu

5. Click on the **Migration(2 of 2)**, and enter the following:

   Migration Thresholds Low   70
   Migration Thresholds High   90

6. Click on the **Apply** push button

ADMIN

You can also use the UPDATE STGPOOL command to perform the preceding task.

Step 8: Stop the Server

You can halt the server without warning if an unexpected problem requires you to return control to the operating system. To avoid losing administrative and client node connections, stop the server only after current sessions have completed or been canceled.

> For most tasks in this book, your server must be running. This procedure is explained here only if an unusual situation requires that you stop the server.

To stop the server, do one of the following:

If the server is running as a service, do the following:

1. Double-click on **Services** icon

2. Click on **Stop** button

This shuts down the server immediately and cancels all client sessions, even if activity has not completed.

1. Double-click on **Server**

2. Click on **Server** from the menu bar

3. Click on **Halt** from the pull down menu

ADMIN

This shuts down the server immediately and cancels all client sessions, even if activity has not completed.

Chapter 3. Customizing Your ADSM System

At this point you have installed, set up, and started the ADSM server, and one or more administrative clients and backup-archive clients. This chapter describes customizing ADSM by doing the following:

- Adding tape drives to your ADSM system

- Automating some client operations and administrative commands

The customization described here does not take advantage of all the options available to you as an ADSM administrator. However, it does provide you with a flexible, powerful ADSM system that can meet most customers' needs.

> **!** The ADSM server must be running to perform the tasks described here. Start the server if it is not running (see "Start the Server" on page 9 for the procedure).

Register ADSM Licenses

When you install ADSM for Windows NT, your system is licensed for the base ADSM support, which provides the following:

- An unlimited number of administrative clients
- One DESKTOP backup-archive client
- Removable media devices not in automated libraries

   Note: For current information about supported devices, check with your authorized reseller. You can also call the IBM Information Support Center at 1-800-IBM-3333 and ask for STAR 20, or you can send an E-mail note to askibm-rsvp@info.ibm.com with STAR 20 in the body of the note.

License certificate files are included with your server package in the server installation directory. Each license certificate file licenses the server for a single option or feature. The file names are:

1CLIENT.LIC	1 backup-archive client
10CLIENT.LIC	10 backup-archive clients
50CLIENT.LIC	50 backup-archive clients
NETWORK.LIC	Using a communications method other than Named Pipes
DESKTOP.LIC	Desktop clients (OS/2, Apple, Novell NetWare, and DOS)
UNIX.LIC	Unix clients
SPACEMAN.LIC	Space management feature
DEVM2TO3.LIC	Upgrade from Device Module 2 to 3
DEVM3TO4.LIC	Upgrade from Device Module 3 to 4
DEVMOD3.LIC	Device Module 3
DEVMOD4.LIC	Device Module 4

To register any ADSM licenses beyond the base support, issue the REGISTER
LICENSE command to read in a license certificate file. For example, to read in the files
that license a total of 50 clients, add the UNIX environment to your base DESKTOP
environment, and allow attaching devices in Device Support Module 3, issue the
following commands:
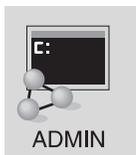
```
register license file(50clients.lic)

register license file(unix.lic)

register license file(network.lic)

register license file(devmod3.lic)
```

ADMIN

Register Additional Administrators

If you have installed any additional administrative clients, you should register them and
grant an authority level to each. For details about the other authority levels, see the
ADSM Administrator's Guide.

> The name SERVER_CONSOLE cannot be used by another
> administrator.

For example, register an administrator with a user ID of MARK and a password of
MISSOURI and grant system authority.

1. Double click on **Administrators**

2. Click on **Edit** from the menu bar

3. Click on **Add** from the pull-down menu

4. In the **Administrator name** field, enter MARK.

5. Click on the **Privilege** tab, and in the **Privilege class** field, click on
   the SYSTEM radio button

ADMIN

6. Click on the **Password** tab, and in the **Password** and **Reenter
   password** fields, enter MISSOURI

7. Click on the **Add** push button

You can also use the REGISTER ADMIN and GRANT AUTHORITY
commands to perform the preceding task.

Change Administrative Client Passwords

To change the password of an administrative client (for example, to TUCSON), do the
following:

1. Double click on **Administrators**

2. From the list of administrator names, click on ADMIN

3. Click on **Selected**

4. Click on **Open as properties** from the pull-down menu

5. Click on the **Password** tab

6. Enter the following:

New password	TUCSON
Reenter new password	TUCSON

7. Click on the **Apply** push button

You can also use the UPDATE ADMIN command to perform the preceding task.

ADMIN

Register Additional Backup-Archive Client Nodes

You are now ready to register any additional backup-archive clients that you have installed. For example, register a node named MERCEDES with the password MONTANA.

1. Double click on **Nodes**

2. Click on **Edit** from the menu bar

3. Click on **Add** from the pull down menu

4. Enter the following:

Node name	MERCEDES

5. Click on the **Password** tab, and enter the following:

Password	MONTANA
Reenter password	MONTANA

6. Click on the **Add** push button

You can also use the REGISTER NODE command to perform the preceding task.

ADMIN

Create New Policies

The ADSM default policy specifies how ADSM manages client files, for example:

- ADSM keeps up to two backup versions of active client files (files existing on clients).

- ADSM backs up client data directly to BACKUPPOOL, a disk storage pool.

You may want to keep the last ten backup versions or back up client data directly to tape. You can modify the default policy or create a new policy. For example, you could create a new management class within the default policy domain. Then the client

include-exclude lists could be updated for each user wanting to use the new management class. On the other hand, you could create a new policy domain and assign clients to that domain. This section describes the second approach.

> **!** **Storing Directly to Tape**
> Be aware that it is usually preferable to store client files to disk and let server migration move the data to tape. When you store directly to tape, the number of available tape drives limits the number of concurrent client sessions that can do meaningful work.

To create a new policy, you can start by copying the policy domain, STANDARD. This operation also copies the associated policy set, management class, and copy groups. You then assign clients to the new domain.

1 Copy the default policy domain, STANDARD, to the new policy domain, NEWDOMAIN.

ADMIN

1. Double click on **Policy Domains**

2. Click on the STANDARD policy domain

3. Click on **Edit** from the menu bar

4. Click on **Add** from the pull-down menu

5. Enter the following:

   Policy domain   NEWDOMAIN

6. Ensure that the **Copy all policy sets, management classes, and copy groups** check box is selected

7. Click on the **Add** push button

You can also use the COPY DOMAIN command to perform the preceding task.

This operation copies the policy domain, and all associated policy sets, management classes, and copy groups. Within the policy domain named NEWDOMAIN and the policy set named STANDARD, you have:

- Management class named STANDARD
- Backup copy group named STANDARD
- Archive copy group named STANDARD

For our example, you update only the backup copy group.

2 Update the backup copy group by specifying that ten versions of backed up files are to be kept.

1. Expand **Policy Domains**

2. Double click on **Backup Copy Groups**

3. Select the copy group from policy domain NEWDOMAIN and policy set STANDARD (you cannot update a copy group from policy set ACTIVE)

4. Click on **File** from the menu bar

5. Click on **Open as properties** from the pull-down menu

6. Click on the **Copy control (2 of 2)** tab, and enter one of the following:

   • To change the number of versions kept:

      Number of backup versions to keep if client data 10

   • To back up directly to tape, autopool for example:

      Destination storage pool AUTOPOOL (from the drop-down list)

7. Click on the **Apply** push button

8. Close the window

You can also use the UPDATE COPYGROUP command to perform the preceding task.

3 Validate the STANDARD policy set in NEWDOMAIN Validation ensures that the active policy set contains a default management class and reports any copy group definition errors.

1. Expand **Policy Domains**

2. Double click on **Policy Sets**

3. Click on the STANDARD policy set in NEWDOMAIN

4. Click on **File** from the menu bar

5. Click on **Validate** from the pull-down menu (**Policy Sets - Validate** is displayed)

6. Click on the **Validate** push button (**Policy Sets - Validate Results** is displayed)

7. Click on the **OK** push button

You can also use the VALIDATE POLICYSET command to perform the preceding task.

The following conditions result in warning messages during validation:

- A destination storage pool is not a defined storage pool.

- The default management class does not contain a backup or archive copy group.

- The current ACTIVE policy set names a management class that is not defined in the policy set being validated.

- The current ACTIVE policy set contains copy groups that are not defined in the named policy set.

- A management class specifies that a backup version must exist before a file can be migrated from a client node, but the management class does not contain a backup copy group.

4 Activate the STANDARD policy set for NEWDOMAIN policy domain.

ADMIN

1. Expand **Policy Domains**

2. Double click on **Policy Sets**

3. Click on the STANDARD policy set in NEWDOMAIN

4. Click on **File** from the menu bar

5. Click on **Activate** from the pull-down menu (**Policy Sets - Activate** is displayed)

6. Click on the **Activate** push button (**Policy Sets - Activate Results** is displayed)

7. Click on the **OK** push button

To verify that the modified backup copy group is now active, do the following:

8. Double click on **Backup Copy Groups**

9. View copy group STANDARD for policy set ACTIVE in NEWDOMAIN and verify that the number of backup versions to keep if client data exists is set to 10

You can also use the ACTIVATE POLICYSET command to perform the preceding task.

5 Assign client nodes to the NEWDOMAIN policy domain by either updating existing client nodes or registering new nodes. For example, to update client node CHLOE, do either of the following:

1. Double click on **Nodes**

2. Click on CHLOE from the list of node names

3. Click on **File** from the menu bar

4. Click on **Open as properties** from the pull-down menu

5. Select NEWDOMAIN in the **Policy domain** field from the drop-down list

6. Click on the **Apply** push button

7. Close the window

You can also use the UPDATE NODE command to perform the preceding task.

ADSM will now use the policy in NEWDOMAIN to manage files for client node CHLOE. If you defined a policy that backs up directly to tape, you can now verify it. From a client assigned to the new domain, follow the steps described in "Step 3: Verify the ADSM Installation" on page 10.

Invoke Expiration Processing

When the retention period of a file in server storage has passed, that file is marked as expired. However, expired files are not deleted from server storage until expiration processing occurs. You can invoke expiration processing in a number of ways:

- The expiration interval server option can specify that processing be invoked at a set interval from the time the server is started. The default is a 24 hour interval.

- You can issue the EXPIRE INVENTORY command manually whenever you wish (see the following example)

- You can schedule the EXPIRE INVENTORY command to be issued regularly at a specific time (see "Automate Administrative Commands" on page 33). If you schedule expiration processing, set the expiration interval server option to 0. You can set the server options by using the Server Options utility on the ADSM Server Utilities.

```
expire inventory
```

Create a Client Include-Exclude List

To make the best use of ADSM storage and policies, each client can exclude certain files from some policy operations and include certain files in other operations. A client includes and excludes files with statements in an include-exclude list (called an include-exclude file on UNIX clients). For example, clients come with a default

include-exclude list that exclude system files that, if recovered, could corrupt the operating system. Consult the appropriate ADSM client user's guide for details.

To help clients get started, you can supply them with additional include-exclude statements tailored to your installation. For example:

- A user may want all files with the extensions **sct** or **drw** in the **spec** directory included for backup services but all other files in that directory excluded. The user could add the following statements to the include-exclude file:

```
exclude d:\eng\spec\*.*
include d:\eng\spec\*.drw
include d:\eng\spec\*.sct
```

> An include-exclude list is read from the bottom up until a match is found. In the preceding example, no match is found on either include statement for a file in the d:\eng\spec directory with an extension of **obj**. The exclude statement is read, a match is found, and the file is excluded.

- For a file or group of files, you can also override the default management class.

```
exclude d:\eng\spec\*.*
include d:\eng\spec\*.drw monthly
include d:\eng\spec\*.sct
```

In this example,

– Files in the **spec** directory with an **sct** extension are bound to the default management class.

– Files with the **drw** extension are bound to the management class **monthly**.

– All other files in the **spec** directory are excluded from backup or archive.

Automate Client Operations

You can automate client operations such as backup to ensure that the clients' data is protected by ADSM. You can automate the following on most clients:

- ADSM backup, archive, restore, and retrieve operations
- Operating system commands
- Macros (files that can contain operating system and ADSM commands)

ADSM for Windows NT supplies you with two default schedules, one for daily incremental backups of client data (DAILY_INCR) and one for weekly backups (**weekly_incr**). In this section, you are guided through the steps necessary to update the daily backup schedule, to define an archive schedule, and to associate those schedules with a group of backup-archive clients. The example assumes that you have registered three client nodes named BILL, MARK, and MERCEDES.

1 The following example changes the time that the schedule begins the backups to 11:00 p.m. and associates the schedule with clients BILL, MARK, and MERCEDES. The schedule, named DAILY_INCR, is for the STANDARD policy domain. The STANDARD policy domain is the ADSM default policy and specifies backup to the disk storage pool BACKUPPOOL.

ADMIN

1. Expand **Central Scheduler**

2. Double click on **Backup/Archive Schedules**

3. Click on the schedule DAILY_INCR

4. Click on **File** from the menu bar

5. Click on **Open as properties** from the pull-down menu

6. Enter the following:

   Policy domain   STANDARD (from the drop-down list)
   Schedule name   DAILY_INCR

7. Associate nodes by doing the following:

   a. Click on the **Add nodes** push button, which displays the **Backup-Archive Schedule - Add Node Associations** window
   b. Click on the **Include** push button
   c. Click on BILL, MARK, and MERCEDES
   d. Click on the **Add** push button

8. Click on the **Timing** tab, and specify the **Start time** as 23:00

9. Click on the **Apply** push button

You can also use the UPDATE SCHEDULE and DEFINE ASSOCIATION commands to perform the preceding task.

2 Define a schedule for an archive operation and associate the schedule with clients BILL, MARK, and MERCEDES. The schedule, named ARC, is for the STANDARD policy domain. The STANDARD policy domain is the ADSM default policy and specifies archive to the disk storage pool ARCHIVEPOOL. This schedule calls for a schedule window that:

- Applies to all files in the c: and d: file spaces
- Begins on the date the schedule is defined (the default) at 11:00 p.m.
- Lasts for 1 hour (the default)
- Is repeated every Friday
- Stays in effect indefinitely (the default)

1. Expand **Central Scheduler**

2. Double click on **Backup/Archive Schedules**

3. Click on **Edit** from the menu bar

4. Click on **Add** from the pull-down menu

5. Enter the following:

   Policy domain   STANDARD (from the drop-down list)
   Schedule name   ARC

6. Associate nodes by doing the following:

   a. Click on the **Add nodes** push button which displays the **Backup-Archive Schedule - Add Node Associations** window
   b. Click on the **Include** push button
   c. Click on BILL, MARK, and MERCEDES
   d. Click on the **Add** push button

7. Click on the **Operation** tab, and do the following:

   a. Click on ARCHIVE from the **Action** drop-down list
   b. In the **Objects** field, specify `c:\* d:\*`

8. Click on the **Timing** tab, and enter the following:

   Start time   23:00
   Period   1
   Period units   WEEKS (from the drop-down list)
   Day of week   FRIDAY (from the drop-down list)

9. Click on the **Add** push button

You can also use the DEFINE SCHEDULE and DEFINE ASSOCIATION commands to perform the preceding task.

ADMIN

3 Start the client scheduler. For the schedules to become active for a workstation, a user must start the scheduler from the node.

END USER

```
dsmc schedule
```

To help ensure that the scheduler is running on the clients, clients can include the SCHEDULE command in their system start-up commands.

The include-exclude list (file on UNIX clients) on each client also affects which files are backed up or archived by the two schedules defined in the preceding steps. For example, if a file is excluded from backup with an EXCLUDE statement, the file will not be backed up when the schedule runs.

4 Because the schedule is to run daily, you can verify that the automation is working as it should on the day after you define the schedule and associate it with clients. If the schedule has run successfully, the status will be *Completed*.

ADMIN

1. Expand the **Central Scheduler** object

2. Double click on **Backup/Archive Events**

3. Look for the schedule and check its status

You can also use the QUERY EVENT command to perform the preceding task.

You can limit the query of events to display only schedules that did not run as expected. On the GUI, click on **View**, click on **Include** from the pull-down menu, and check the **Exceptions only** check box.

Schedules that did not complete successfully have a status of *Failed*, *Missed*, or *Severed*.

5 Check the results of the schedule on one of the clients that was associated with that schedule. For most clients, including Windows 32-bit, information about what happens when a schedule runs is stored in the file dsmsched.log.

Automate Administrative Commands

You can automate ADSM administrative tasks by scheduling administrative commands. For example, you can automate commands that need to be run regularly or that require significant processing or other system resources. In this section, you define a schedule to run expiration processing.

1 Define a schedule named EXPPROC that invokes expiration processing every night at 2:00 a.m.

This schedule calls for a schedule window that:

- Begins on the date the schedule is defined (the default) at 2:00 a.m .
- Lasts for 1 hour (the default)
- Is repeated every day
- Takes effect immediately
- Stays in effect indefinitely (the default)

ADMIN

1. Expand the **Central Scheduler** object

2. Double click on **Administrative Command Schedules**

3. Click on **Edit** from the menu bar

4. Click on **Add** from the pull-down menu

5. Enter the following:

   Schedule name   EXPPROC
   Administrative Command   expire inventory

6. Click on the **Timing** tab, and specify the **Start time** as 02:00

7. Click on the **Add** push button

You can also use the DEFINE SCHEDULE command to perform the preceding task.

2 Because the EXPPROC schedule is to run daily, you can verify that the automation is working as it should on the day after you define the schedule. If the schedule has run successfully, the status will be *Completed*.

ADMIN

1. Expand the **Central Scheduler** object

2. Click on **Administrative Command Events**

3. Check the status of the EXPPROC schedule

You can also use the QUERY EVENT command to perform the preceding task.

Protect ADSM and Client Data

The database, recovery log, and storage pools are critical to the operation of the server. If the database or recovery log is unusable, the entire server is unavailable. If a database is lost and cannot be recovered, the backup, archive, and space-managed data for that server is lost. If a storage pool volume is lost and cannot be recovered, the data on the volume is also lost.

ADSM provides a number of ways to protect your data, including backing up your storage pools and database. For example, you can define schedules so that the following operations occur:

- After the initial full backup of your storage pools, incremental storage pool backups are done nightly.

- Full database backups are done weekly.

- Incremental database backups are done nightly.

See the *ADSM Administrator's Guide* for details.

Appendix A. Using ADSM Administrative Interfaces

This chapter describes how to use the following ADSM administrative interfaces:

Using the ADSM Server Utilities Interface

The following is a brief description of the ADSM Server Utilities (see Figure 11).
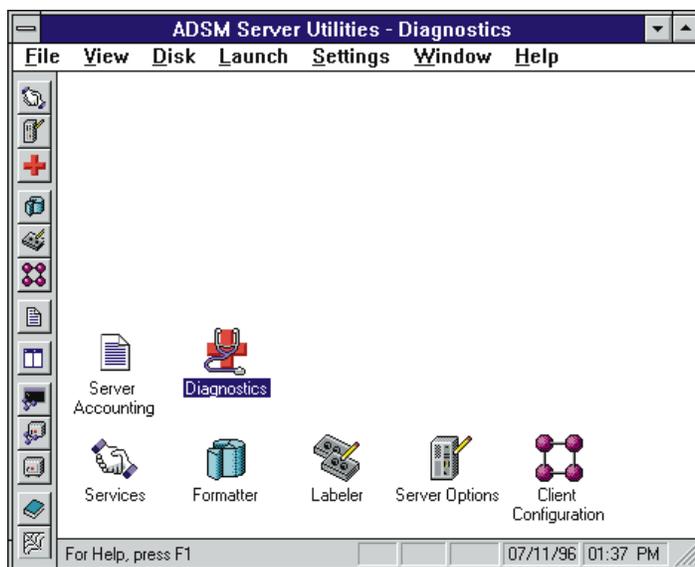
Figure 11. ADSM Server Utilities

Formatter

Helps estimate the size of and create disk volumes for storage pools, the database, and the recovery log.

Services

Helps start the ADSM server, scheduler, and SCSI device driver as Windows NT services.

Labeler

Labels tapes for ADSM

Server Options

Helps update server options. See "Setting Server Communications" on page 41 for communications options information and see *ADSM Administrator's Reference* for other server options.

Client Configuration
　　Helps create create client options and a a file that represents the local
　　ADSM server. This file can be read by a client configuration utility running
　　on remote clients.

Server Accounting
　　Displays a file which includes accounting information.

Diagnostics
　　Displays the installed paths for ADSM components and information about
　　attached devices.

Using the Graphical User Interface

You can manage the server from a graphical user interface. For instructions on starting
the Windows 32-bit GUI, see "Step 6: Start and End an Administrative Client Session"
on page 12.

You can use the GUI to do the following tasks:

* Define, change, or delete an ADSM object
* View ADSM server and database information
* Invoke other ADSM functions, for example, activating a policy set

Defining an Object

The following example shows how to define an ADSM client node:

1 Double click on the **Nodes** icon.

2 Click on **Edit** from the menu bar and then **Add** from the pull down menu.

　　The **Node - Add** notebook appears.

3 The **Node - Add** notebook contains notebook tabs. Click on the tabs to page
　　through the notebook.

4 Define the node by entering data, clicking on radio buttons, selecting check
　　boxes, and selecting from drop-down lists.

5 When you have completed your definition, click on the **Add** push button.

Changing an Object

The following example shows how to change information for an ADSM client node:

1 Double click on the **Nodes** icon.

　　The **Nodes - Details** screen appears.

2 Click on a node.

3 Click on **File** from the menu bar and then **Open as properties** from the pull down menu.

The **Properties** notebook appears.

4 The **Properties** notebook contains notebook tabs. Click on the tabs to page through the notebook.

5 Make your changes by entering data, clicking on radio buttons, selecting check boxes, and selecting from drop-down lists.

6 When you have made all your changes, click on the **Apply** push button.

Deleting an Object

The following example shows how to delete an ADSM client node:

1 Double click on the **Nodes** icon.

2 Click on the node or nodes that you want to delete.

3 Click on **Edit** from the menu bar and then **Delete** from the pull down menu.

The **Nodes - Delete** notebook appears.

4 Verify the node name or names and click on the **Delete** push button.

Using Online Help

The administrative client GUI provides the following types of online help:

Task help	Step-by-step information on how to complete the task associated with a selected object.
General help	An overview of the function of a window. For example, general help for the Nodes—Icons window explains how to use the window, lists the tasks you can complete with the Nodes object, and displays links where you can obtain more information.
Contextual help	Specific information for a selectable object, menu choice, notebook tab, field, and control or push button in a window.
Help Topics	Starts the online help system and displays the Help table of contents.
Using help	Information about the types of help.
Hypertext links	Links to more information about words that appear in a different color from the rest of the help text.
Find	Under Services on a help window, you can search for occurrences of a text string in the online help.
View Books	You can view online ADSM books.

Using the Administrative Command-Line Interface

The administrative command-line client lets administrators control and monitor the server through administrative commands. For more information on using the command-line interface, see *ADSM Administrator's Reference*.

On Windows NT, the administrative client invokes ADSM by using the `DSMADMC` command. This section shows Windows NT examples. For more details and examples of other platform invocations, see *ADSM Administrator's Reference*.

Using Interactive Mode

To enter a series of administrative commands, start the administrative client in interactive mode by entering:

```
dsmadmc
```

You are then prompted to enter your user ID and password. If you do not want to be prompted for that information, you can enter your user ID and password with the DSMADMC command by using the -ID and -PASSWORD options. For example, enter:

```
dsmadmc -id=smith -password=secret
```

To end an administrative client session in interactive mode, use the QUIT command:

```
quit
```

Appendix B. Server and Client Communications Options

This appendix assists you in setting your Windows NT communications.

Setting Server Communications

Use the ADSM Server Utilities to help you set up and customize your ADSM system. You can view or modify the server communication parameters from the Server Options utility.

Communication settings can also be configured for your computer using the network control panel (see "Procedure to Update NETBIOS from Windows NT Control Panel" on page 44 for an example). This may be necessary to enable a particular protocol or to adjust system parameters.

By default, the server listens for incoming ADSM client connections on TCP/IP and Named Pipes. If you start the server console and see warning messages that a particular protocol could not be used by the server, either the protocol is not installed on your machine or the ADSM settings do not match the Windows NT protocol settings.

You can also use the ADSM Server Options utility to get information that clients need to communicate with the server. For a client to use a protocol that is enabled on the server, the client options file must contain corresponding values for communication options. You can view these values for each protocol from the Server Options utility.

! The ADSM Server Utilities allow you to specify communication methods described in this book. For details about other options, such as date and time formats, and buffer sizes, see *ADSM Administrator's Reference* or the online help.

This section describes setting server options *before* you start the server. When you start the server, the new options go into effect. If you modify any server options after starting the server, you must stop and restart the server to activate the updated options.

TCP/IP Communication Method

Figure 12 on page 42 displays an example of a TCP/IP communication method setting from the ADSM Utilities. Each of the communication methods has a screen similar to the one shown here. Click on the **TCP/IP Parameters** button for server and client parameters. You can enable or disable the communications connection with the check box on this screen.
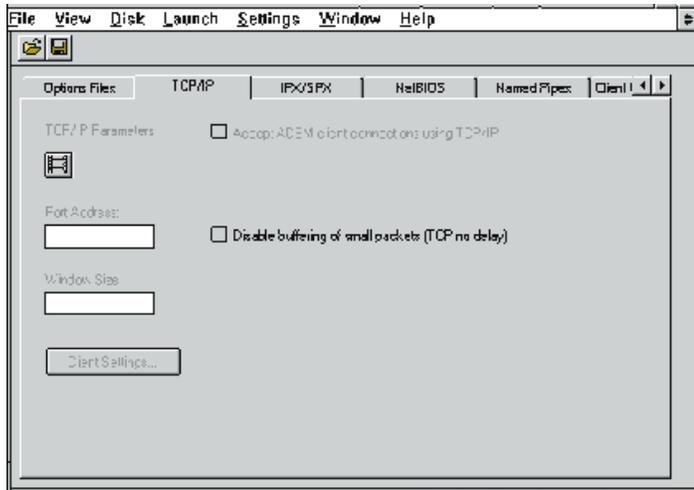
Figure 12. TCP/IP Server Option Screen

```
COMMMETHOD              TCPip
   TCPPort              1500
   TCPWindowsize        8
   TCPNODELAY           NO
```

Figure 13. Example of TCP/IP Communication Method Options

TCPPort
Specifies a port address of a server when using TCP/IP. The default value is 1500.
You can specify a value from 1024 to 32767.

TCPWindowsize
Specifies the amount of data to send or receive before TCP/IP exchanges
acknowledgments with the client node. The actual window size used in a session is the
minimum size of the server and client window sizes. Larger window sizes use
additional memory but may improve performance.

TCP/IP windowsize is specified in kilobytes. You can specify a value from 0 to 2048. 0
indicates that the default window size set for Windows NT should be used. Values
from 1 to 2048 indicate that the window size ranges from 1KB to 2MB.

The default window size is 8.

TCPNODELAY
Specifies whether the server should send small amounts of data or allow TCP/IP to
buffer the data. Disallowing buffering may improve throughput but more packets are
sent over the network.

The default value is NO.

YES   Specifies that TCP/IP should not buffer small ADSM messages

NO   Specifies that TCP/IP should buffer small ADSM messages

No additional server configuration is required for TCP/IP support of multiple adapters. Each adapter must have a separate TCP/IP address. Clients are given one address for their client options file. The server accepts sessions from each of the adapters. The session limit is placed on the accumulated count from all of the adapters.

NETBIOS Communication Method

On Windows NT, NetBIOS can be used to communicate over TCP/IP, IPX/SPX, or NetBEUI. Two machines can communicate if their communications method and underlying transport match. If you move one transport enabler, you see each binding assigned to a LANADAPTER. This is repeated if you have multiple physical network adapters.

Some clients communicate using NetBIOS over the NetBEUI adapter. For example, to have an OS/2 client back up to an ADSM server on Windows NT, the LANADAPTER statement in the server options file must match the entry in the system control panel that has the Nbwnt->Nbf bindings.

Figure 14 displays an example of a Network Basic/Input Output (NetBIOS) communication method setting.

```
COMMMETHOD              NETBIOS
   NETBiosname          NTSERVER1
   LANAdapter           2
   NETBIOSSessions      16
   NETBIOSBuffersize    32
```

Figure 14. Example of NETBIOS Communication Method Options

NETBiosname
Specifies the name to be used for communications on the network. This name must be unique across the network that includes the LAN requestors and other NETBIOS applications. This name can be a maximum of 16 characters and is case-sensitive. The first character cannot be an asterisk (*), and the first three characters cannot be the letters IBM. The defaults are NTSERVER1, NTSERVER2, NTSERVER3, and NTSERVER4 (depending on the number of LANAdapter statements you have). For example, if you have LANADAPTER 0 and LANADAPTER 1, the defaults are NTSERVER1 and NTSERVER2.

LANAdapter
Specifies a network adapter number on which the server communicates when the NETBIOS communication method is specified. The default value for the first LANAdapter statement is 2. You must specify a LANAdapter option for each adapter card you are using with NETBIOS. The NETbiosname statement then specifies the

server names, one per adapter on which the ADSM server accepts sessions. Clients use the NETBIOSservername statement to name the corresponding name on the server.

ADSM requires NETBEUI protocol support to communicate between Windows NT and OS/2.

NETBIOSSessions
Specifies the number of NETBIOS sessions that are reserved for this process per adapter. The minimum number of sessions is 1 and the maximum number of sessions is 255.

This is independent of the MAXSESSION server option, other than MAXSESSION is the default number of NETBIOS sessions. THE MAXSESSION limit overrides the NETBIOSSESSIONS limit for allowing a session to start.

NETBIOSBuffersize
Specifies the size you want to use in kilobytes for the NETBIOS communications buffer. This parameter is optional and the default buffer size is 32KB, with a minimum of 1KB and a maximum of 32KB. A larger buffer size improves communication performance but requires more memory.

Procedure to Update NETBIOS from Windows NT Control Panel

You can update communications from the Windows NT program by following this procedure.

1 Return to the Windows NT screen.

2 Double-click on the **Control Panel** icon.

3 Double-click on the **Network** icon.

4 At the **Network Settings** window, select **NetBEUI Protocol** and click on **Bindings**. See Figure 15 on page 45.

5 Return to the **Network Settings** screen.

6 Double-click on **NetBIOS Interface**.

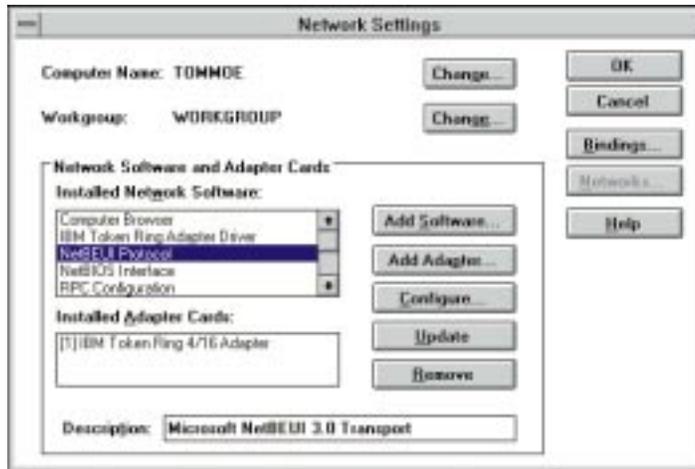**7** Ensure that Nbf is assigned to LANAdapter 2.

Figure 15. Network Settings for NetBIOS

NetWare IPX/SPX Communication Method

Figure 16 displays an example of Internetwork Packet Exchange/Sequenced Packet Exchange (IPX/SPX) communication method setting.

```
COMMMETHOD                IPXSPX
     IPXSocket            8522
     IPXBuffersize        32
```

Figure 16. Example of IPX/SPX Communication Method Options

IPXSocket

Specifies the socket number (hex value) on which the server SPX communication driver is to wait. The default is 8522. You can specify a value from 0 to 32767. Contact your system administrator for this value.

Note: The default socket number is registered with Novell.

IPXBuffersize

Specifies the size you want to use in kilobytes for the IPX/SPX communications buffer. This parameter is optional and the default buffer size is 32KB, with a minimum of 1KB and a maximum of 32KB. A larger buffer size improves communication performance, but requires more memory.

Note: If you are updating IPX/SPX communication method by using the procedure in "Procedure to Update NetWare IPX/SPX from Windows NT Control Panel" on page 46, you must set **Auto Frame Type Detection**.

Procedure to Update NetWare IPX/SPX from Windows NT Control Panel

You can update communications from the Windows NT program by following this procedure.

1 Return to the Windows NT screen.

2 Double-click on the **Control Panel** icon.

3 Double-click on the **Network** icon.

4 At the **Network Settings** window, select **NWLink IPX/SPX Compatible Transport** and click on **Configure**.

5 Ensure that the frame type is **Auto Detected**.

6 Click on **OK**.

Named Pipes Communication Method

Figure 17 displays an example of a Named Pipes communication method setting. This communication method is ideal when running the server and client on the same Windows NT machine because Named Pipes support is internal to the Windows NT base system. Named Pipes require no special configuration.

```
COMMMETHOD                        NAMEdpipe
     NAMEDPIPENAME                \\.\PIPE\ADSMPIPE
```

Figure 17. Example of Named Pipes Communication Method Options

NAMEDPIPENAME

Specifies the Named Pipes name. The default is \\.PIPE\ADSMPIPE.

Setting Up the Windows 32-bit Client

This chapter describes setting up the Windows administrative client. It also provides some information on setting client communications options.

> **!** For information about installing administrative clients and backup-archive clients other than Windows 32-bit clients, see *ADSM Installing the Clients*.

The ADSM Windows 32-bit administrative and backup-archive clients are installed when you install the ADSM package. Each of these components has a corresponding client options file which you can edit by clicking on the **Backup-archive Client Options** or the **Administrative Client Options** icons respectively. You can use the ADSM Utilities Client Configuration icon (see Figure 11 on page 37) or the Administrative Client Options icon (see Figure 10 on page 7) to change client options.

Setting ADSM for Windows 32-bit Client Communications Options

The client user options file identifies the server you can connect to and specifies the communication method and associated options.

ADSM for Windows NT clients use the communication methods shown in Table 2 to communicate with the server.

Table 2. Supported Communication Methods for Administrative Clients

Communication Method	Supported Clients	Servers
TCPIP	AIX, AT&T UNIX, Bull DPX/2, Digital UNIX, DEC ULTRIX, DOS, HP-UX, NEW EWS-UX/V, MVS TSO, OpenEdition MVS, OS/2, SCO UNIX 386/Open Desktop, Sequent PTX, Siemens Nixdorf SINIX, Silicon Graphics IRIX, Sun Microsystems SunOS or Sun Solaris, VM CMS, Windows, Windows 32-bit	AIX, AS/400, HP-UX, MVS, OS/2, Sun Microsystems Sun Solaris, VM, Windows NT
NETBIOS	DOS, OS/2, Windows, Windows 32-bit	AIX, OS/2, Windows NT
IPXSPX	DOS, Novell, Windows, Windows 32-bit	AIX, AS/400. OS/2, Windows NT
NAMED PIPES See Note	OS/2, Windows 32-bit	OS/2, Windows NT

Note:

- An OS/2 server can talk with OS/2 clients using Named Pipes.

- Windows NT servers can talk with Windows 32-bit clients using Named Pipes. The Windows 32-bit client can be on a different machine than the server.

See *ADSM Installing the Clients* for additional information.

The ADSM client uses TCP/IP, IPX/SPX, NETBIOS, or Named Pipes to communicate with the Windows NT server.

TCP/IP Communication Method

Figure 18 displays an example of a TCP/IP communication method setting.

```
COMMMETHOD               TCPip
   TCPServeraddress      localhost
   TCPPort               1500
   TCPBuffsize           31
   TCPWindowsize         8
```

Figure 18. TCP/IP Communication Method Options for a Windows 32-bit Client

TCPServeraddress

Specifies the Internet server address for communication using TCP/IP. This address can be a TCP/IP Internet Domain Name (for example, gldvm.end.ibm.com), or a dot address (for example, 9.103.21.200).

TCPPort

Specifies a port address of a server when using TCP/IP.

The default value is 1500. You can specify a value from 1000 to 32767. Contact your system administrator for this value.

TCPBuffsize

Specifies the size of the internal TCP communication buffer, which is used to transfer data between the client node and the server. A larger buffer size can improve communication performance but requires more memory.

Specify the size of the buffer in kilobytes. The default value is 31KB. You can specify a value from 1 to 32KB.

TCPWindowsize

Specifies the size of the TCP window for your client. A larger window size can improve communication performance, but uses more memory.

You can specify a value from 1 to 24KB. The default is 8KB.

Options for NETBIOS Communication Method

Figure 19 shows an example of a NETBIOS communication method setting used by an Windows 32-bit client.

```
COMMMETHOD               NETBIOS
   NETBIOSServername     NTSERVER1
   NETBIOSName           CLIENT
   NETBIOSBuffersize     16
   NETBIOSTimeout        60
   LANAdapter            2
```

Figure 19. Example of NETBIOS Communication Method Options for the Client

NETBIOSServername

Specifies the NETBIOS name for an ADSM server. The default for this option is **NTSERVER1**.

NETBIOSName

Specifies the NETBIOS name for your workstation. The default for this option is **ADSM**concatenated with a time-of-day value. Normally, the default value should be used because it causes a unique name to be generated for each session.

NETBIOSBuffersize

Specifies the size you want to use, in kilobytes, for the NETBIOS communications buffer. The default is 16KB and the maximum is 32KB. The recommended size is 24KB.

NETBIOSTimeout

Specifies the number of seconds that must elapse before a timeout occurs for a NETBIOS send or receive. The default is 60 seconds.

LANAdapter

Specifies the communication adapter number. This number must be set to 2.

Options for IPX/SPX Communication Method

Figure 20 shows an example of an IPX/SPX communication method setting used by a Windows 32-bit client.

```
COMMMETHOD               IPXSPX
   IPXServeraddress      000000208005a30bee98
   IPXSOcket             8522
   IPXBuffersize         16
```

Figure 20. Example of IPX/SPX Communication Method Options for the Administrative Client

IPXServeraddress

Specifies a 20 digit hexadecimal number address for an ADSM server. The first 8 hexadecimal digits are for the ADSM server IPX network address. The next 12 digits are for the ADSM server node IPX address. This address is obtained by running the **getipxad** utility with the socket number on the system to be used as a server. This file can be found in the **c:\win32app\ibm\adsm\utils** directory.

This screen shows the **getipxad** command and the results of the command:

```
C:\win32app\ibm\adsm\utils>getipxad 8522

Network number: 00000002
Node number:    10005AC44B92
Socket number:  8522

To communicate with this server, put the following in your
ADSM client options file:

ipxsocket         8522
ipxserveraddress  0000000210005AC44B92
```

If you do not specify an address, ADSM assumes the server resides on your client
workstation.

IPXSOcket

Specifies the socket number on which the server SPX communication driver is to wait
for requests. The socket number is a hex value on which the server listens. The
default is 8522.

Note: The socket number is registered with Novell.

IPXBuffersize

Specifies the size you want to use, in kilobytes, for the IPXSPX communications buffer.
This parameter is optional and the default buffer size is 32KB, with a minimum of 1KB
and a maximum of 32KB.

Appendix C. Where to Look for More Information

This section shows you where to look for more information.

ADSM Publications

The following table lists the ADSM publications available in hardcopy and CD-ROM.

Short Title	Publication Title	Order Number
Hardcopy Publications		
ADSM General Information	*ADSTAR Distributed Storage Manager: General Information*	GH35-0131
ADSM Messages	*ADSTAR Distributed Storage Manager: Messages*	SH35-0133
ADSTAR Distributed Storage Manager for Windows NT: Administrator's Guide	*ADSTAR Distributed Storage Manager for Windows NT: Administrator's Guide*	GC35-0236
ADSTAR Distributed Storage Manager for Windows NT: Administrator's Reference	*ADSTAR Distributed Storage Manager for Windows NT: Administrator's Reference*	GC35-0237
ADSTAR Distributed Storage Manager for Windows NT: License Information	*ADSTAR Distributed Storage Manager for Windows NT: License Information*	GC35-0234
ADSM Using the UNIX HSM Clients	*ADSTAR Distributed Storage Manager: Using the UNIX Hierarchical Storage Management Clients*	SH26-4030
ADSM V2 Using the Apple Macintosh Backup-Archive Client	*ADSTAR Distributed Storage Manager Version 2: Using the Apple Macintosh Backup-Archive Client*	SH26-4051
ADSM Using the UNIX Backup-Archive Clients	*ADSTAR Distributed Storage Manager Version 2: Using the UNIX Backup-Archive Clients*	SH26-4052
ADSM V2 Using the OS/2 Backup-Archive Client	*ADSTAR Distributed Storage Manager Version 2: Using the OS/2 Backup-Archive Client*	SH26-4053
ADSM V2 Using the DOS Backup-Archive Client	*ADSTAR Distributed Storage Manager Version 2: Using the DOS Backup-Archive Client*	SH26-4054
ADSM V2 Using the Novell NetWare Backup-Archive Client	*ADSTAR Distributed Storage Manager Version 2: Using the Novell NetWare Backup-Archive Client*	SH26-4055
ADSM V2 Using the Microsoft Windows Backup-Archive Clients	*ADSTAR Distributed Storage Manager Version 2: Using the Microsoft Windows Backup-Archive Clients*	SH26-4056

Short Title	Publication Title	Order Number
ADSM Using the Lotus Notes Backup Agent	*ADSTAR Distributed Storage Manager: Using the Lotus Notes Backup Agent*	SH26-4047
ADSM Installing the Clients	*ADSTAR Distributed Storage Manager: Installing the Clients*	SH26-4049
ADSM Client Reference Cards	*ADSTAR Distributed Storage Manager: Client Reference Cards*	SX26-6013
ADSM Using the Application Programming Interface	*ADSTAR Distributed Storage Manager: Using the Application Programming Interface*	SH26-4002
CD-ROM		
ADSM Online Library		SK2T-1893

IBM Redbooks

The International Technical Support Center (ITSC) publishes redbooks, books on specialized topics such as using ADSM to back up databases. You can order publications through your IBM representative or the IBM branch office serving your locality. You can also search for and order books of interest to you by visiting the IBM Redbooks home page on the World Wide Web at this address:

> http://www.redbooks.ibm.com/redbooks

Software Developer's Program

The IBM Storage Systems Division (SSD) Software Developer's Program provides a range of services to software developers who want to use the ADSM application program interface (API). Information about the SSD Software Developer's Program is available in:

- IBMSTORAGE forum on CompuServe
- SSD Software Developer's Program Information Package

To obtain the Software Developer's Program Information Package:

1. Call 800-4-IBMSSD (800-442-6773). Outside the U.S.A., call 408-256-0000.
2. Listen for the Storage Systems Division Software Developer's Program prompt.
3. Request the Software Developer's Program Information Package.

Do You Have Comments or Suggestions?

If you have difficulty using this publication or if you have comments and suggestions for improving it, please complete and mail the readers comment form found in the back of this publication. Your comments and suggestions can contribute to the quality and usability of this publication.

Translations

Selected ADSM publications have been translated into languages other than American English. For a complete list of the available translations and their order numbers, see *IBM ADSTAR Distributed Storage Manager General Information*. Contact your IBM representative for more information about the translated publications and whether these translations are available in your country.

Appendix D. Notices

References in this publication to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any of the intellectual property rights of IBM may be used instead of the IBM product, program, or service. The evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM, are the responsibility of the user.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to the IBM Director of Licensing, IBM Corporation, 500 Columbus Avenue, Thornwood, New York 10594, U.S.A.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact IBM Corporation, Information Enabling Requests, Dept. M13, 5600 Cottle Road, San Jose CA 95193, U.S.A. Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

Trademarks

The following terms are trademarks of the IBM Corporation in the United States or other countries or both:

ACF/VTAM	DFSMS	Proprinter
AD/Cycle	DFSMS/MVS	PS/2
ADSTAR	DFSMSdss	RACF
Advanced Peer-to-Peer Networking	ESCON	RISC System/6000
AIX	IBM	RS/6000
AIX/6000	IBMLink	SAA
AIXwindows	Language Environment	SP2
Application System/400	Library Reader	System/370
APPN	MVS/DFP	System/390
AS/400	MVS/ESA	Systems Application Architecture
AT	MVS/SP	SystemView
BookManager	MVS/XA	Virtual Machine/Enterprise Systems Architecture
C/370	OpenEdition	Virtual Machine/Extended Architecture
CICS	Operating System/2	VM/ESA
Common User Access	Operating System/400	VM/XA
CUA	OS/2	VSE/ESA
DATABASE 2	OS/400	VTAM
DB2/6000	POWERparallel	WIN-OS/2

The following terms are trademarks of other companies:

Trademark	Company	Trademark	Company
Andataco	Andataco Corporation	NetWare	Novell, Inc.
Apple	Apple Computer, Inc.	NFS	Sun Microsystems, Inc.
Attachmate	Attachmate Corporation	Novell	Novell, Inc.
CompuServe	CompuServe, Inc.	Open Desktop	The Santa Cruz Operation, Inc.
dBASE	Borland International, Inc.	OpenWindows	Sun Microsystems, Inc.
DECstation	Digital Equipment Corporation	PARADOX	Borland International, Inc.
DLT	Quantum Corporation	PC/TCP	FTP Software, Inc.
DPX/20	Groupe Bull	PTX	Sequent Computer Systems
Dynatek	Dynatek Automation Systems	SCO	The Santa Cruz Operation, Inc.
DynaText	Electronic Book Technologies, Inc.	Sequent	Sequent Computer Systems
Exabyte	Exabyte Corporation	SINIX	Siemens Nixdorf Information Systems,
Extra!	Attachmate Corporation		Inc.
FOXPRO	Microsoft Corporation	Solaris	Sun Microsystems, Inc.
Hewlett-Packard	Hewlett-Packard Company	Sony	Sony Corporation
HP-UX	Hewlett-Packard Company	SPARC	SPARC International, Inc.
Ice Box	Software International Microsystems	StorageTek	Storage Technology Corporation
iFOR/LS	Gradient Technologies, Inc.	Sun	Sun Microsystems, Inc.
INGRES	ASK Group, Inc.	Sun Microsystems	Sun Microsystems, Inc.
Intel	Intel Corporation	SunOS	Sun Microsystems, Inc.
Iomega	Iomega Corporation	Sun-3	Sun Microsystems, Inc.
IPX/SPX	Novell, Inc.	Sun-4	Sun Microsystems, Inc.
IRIX	Silicon Graphics, Inc.	SureStore	Hewlett-Packard Company
Jetstore	Hewlett-Packard Company	SyQuest	SyQuest Technology, Inc.
Lotus	Lotus Development Corporation	ULTRIX	Digital Equipment Corporation
Lotus Notes	Lotus Development Corporation	WangDAT	Tecmar Technologies, Inc.
Macintosh	Apple Computer, Inc.	Windows 95	Microsoft Corporation
MacTCP	Apple Computer, Inc.	Windows NT	Microsoft Corporation
Motif	Open Software Foundation, Inc.	X Windows	Massachusetts Institute of Technology

C-bus is a trademark of Corollary, Inc.

Microsoft, Windows, and the Windows 95 logo are trademarks or registered trademarks of Microsoft Corporation.

PC Direct is a trademark of Ziff Communications Company and is used by IBM Corporation under license.

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Limited.

Other company, product, and service names, which may be denoted by a double asterisk (**), may be trademarks or service marks of others.

Glossary

The terms in this glossary are defined as they pertain to the ADSM library. If you do not find the term you are looking for, refer to the *IBM Dictionary of Computing*, New York: McGraw-Hill, 1994.

This glossary may include terms and definitions from:

- The *American National Standard Dictionary for Information Systems*, ANSI X3.172-1990, copyright (ANSI). Copies may be purchased from the American National Standards Institute, 11 West 42nd Street, New York 10036.
- The *Information Technology Vocabulary*, developed by Subcommittee 1, Joint Technical Committee 1, of the International Organization for Standardization and the International Electrotechnical Commission (ISO/IEC JTC2/SC1).

A

absolute.   A backup copy group mode value indicating that a file is considered for incremental backup even if the file has not changed since the last backup. See also *mode*. Contrast with *modified*.

access mode.   A storage pool and storage volume attribute that specifies whether data can be written to or read from storage pools or storage volumes. The access mode can be read/write, read-only, or unavailable. Volumes in primary storage pools can also have an access mode of destroyed. Volumes in copy storage pools can also have an access mode of offsite.

accounting facility.   A facility that records statistics about client session activity.

accounting records.   Files that record session resource usage at the end of each client session.

action choice.   A choice in a pull-down menu that causes an action. See also *routing choice*.

activate.   The process of validating the contents of a policy set and copying the policy set to the ACTIVE policy set.

active policy set.   The policy set within a policy domain that contains the most recently activated policy currently in use by all client nodes assigned to that policy domain. See *policy set*.

active version.   The most recent backup copy of a file stored by ADSM. Such a file is exempt from deletion until a backup detects that the user has either replaced the file with a newer version, or has explicitly deleted the file from the workstation. Contrast with *inactive version*.

activity log.   A log that records normal activity messages generated by the server. These messages include information about server and client operations, such as the start time of sessions or device I/O errors. Each message includes a message ID, date and time stamp, and a text description. The number of days to retain messages in the activity log can be specified.

administrative client.   A program that runs on a file server, workstation, or mainframe that allows administrators to control and monitor the server through administrator commands. Contrast with *backup-archive client.*

administrative command schedule.   A database record that describes the planned processing of an administrative command during a specific time period. See also *client schedule*.

administrative privilege class.   A permission granted to an administrator that controls the commands that the administrator can issue. See *system privilege class, analyst privilege class, operator privilege class, policy privilege class or storage privilege class*.

administrative session.   A period of time in which an administrator user ID can communicate with a server to perform administrative tasks. Contrast with *client node session*.

administrator.   A user who has been registered to the server. Administrators can be authorized to one or more of the following administrative privilege classes: system, policy, storage, operator, or analyst. Administrators can use the administrative client to enter server commands and queries in accordance with their privileges.

administrator definition.   Server control information that includes the administrator's name, password, contact information, administrative privilege classes, policy domains and storage pools assigned to an administrator, and whether the administrative ID is locked from the server. An administrator definition can be exported from a source server and imported to a target server at a later date.

ADSM.   ADSTAR Distributed Storage Manager.

ADSM application program interface (API).   A set of functions that applications running on a client platform can call to store, query, and retrieve objects from ADSM storage.

ADSTAR Distributed Storage Manager (ADSM).   A client/server program that provides storage management to customers in a multivendor computer environment.

Advanced Interactive Executive (AIX).   An operating system used in the RISC System/6000 computers. The AIX operating system is IBM's implementation of the UNIX operating system.

Advanced Peer-to-Peer Networking (APPN).   An extension to the LU6.2 peer orientation for end-user services. See *SNA LU6.2* and *Systems Network Architecture*.

Advanced Program-to-Program Communication (APPC).   An implementation of the SNA/SDLC LU6.2 protocol that allows interconnected systems to communicate and share the processing of programs. See *SNA LU6.2*, *Systems Network Architecture*, and *Common Programming Interface Communications*.

AFS.   Andrew file system.

AIX.   Advanced Interactive Executive.

analyst privilege class.   An administrative privilege class that allows an administrator to reset statistics.

Andrew file system (AFS).   A distributed file system developed for UNIX operating systems.

API.   Application program interface.

APPC.   Advanced Program-to-Program Communication.

APPN.   Advanced Peer-to-Peer Networking.

archive.   A function that allows users to copy one or more files to a storage pool for long-term storage. Archive copies may be accompanied by descriptive information and may be retrieved by archive date, by file name, or by description. Contrast with *retrieve*.

archive copy.   A user file that has been archived to an ADSM storage pool.

archive copy group.   A policy object containing attributes that control the generation, destination, and expiration of archive files. An archive copy group belongs to a management class.

ARCHIVEPOOL.   A disk storage pool defined by ADSM at installation. It can be the destination for client files that are archived to the server. See *storage pool*.

archive retention grace period.   The number of days ADSM retains an archive copy when the server is unable to rebind the file to an appropriate management class.

AS/400.   Application System/400.

assigned capacity.   The portion of available space that can be used to store database or recovery log information. See also *available space*.

association.   The relationship between a client node and a client schedule. An association identifies the name of a schedule, the name of the policy domain to which the schedule belongs, and the name of a client node that performs scheduled operations.

audit.   The process of checking for logical inconsistencies between information that the server has and the actual condition of the system. ADSM has processes for auditing volumes, the database, libraries, and licenses. For example, in auditing a volume ADSM checks for inconsistencies between information about backed up or archived files stored in the database and actual data associated with each backup version or archive copy in server storage.

authentication.   The process of checking a user's password before allowing that user access to the server. Authentication can be turned on or off by an administrator with system privilege.

autochanger.   A small multislot tape device that has a mechanism that automatically puts tape cartridges into the tape drive or drives. Also called *medium* or *media changer*, or a *library*.

available space.   The amount of space, in megabytes, that is available to the database and recovery log. This space can be used to extend the capacity of the database or recovery log, or to provide sufficient free space before a volume is deleted from the database or recovery log.

B

background process. A server process that runs in the background, allowing the administrative client to be used for other work.

backup. The process of copying information for safekeeping. ADSM has processes for backing up user files, the database, and storage pools. For example, users can back up one or more files to a storage pool to ensure against loss of data. Contrast with *restore*. See also *database backup series* and *incremental backup*.

backup-archive client. A program that runs on a workstation or file server and provides a means for users to back up, archive, restore, and retrieve files. Contrast with *administrative client*.

backup copy. A user file that has been backed up to an ADSM storage pool.

backup copy group. A policy object containing attributes that control the generation, destination, and expiration of backup files. A backup copy group belongs to a management class.

BACKUPPOOL. A disk storage pool defined by ADSM at installation. It can be the destination for client files that are backed up to the server. See *storage pool*.

backup retention grace period. The number of days ADSM retains a backup version after the server is unable to rebind the file to an appropriate management class.

backup series. See *database backup series*.

backup version. A file, directory, or file space that a user has backed up, which resides in ADSM server storage. There may be more than one backup version of a file in the storage pool, but at most only one is an active backup version. See *active version* and *inactive version*.

binding. The process of associating a file with a management class name. See *rebinding*.

buffer. Storage used to compensate for differences in the data rate flow, when transferring data from one device to another.

buffer pool. Temporary space used by the server to hold database or recovery log pages. See *database buffer pool* and *recovery log buffer pool*.

buffer pool size. The size of an area in memory used to store database or recovery log pages.

bus converter. A device that translates between different Hewlett-Packard internal I/O bus architectures.

C

cache. The process of leaving a duplicate copy on random access media when the server migrates a file to another storage pool in the hierarchy.

cartridge. A sequential storage media that contains magnetic tape in a protective housing. Contrast with *tape reel*.

CARTRIDGE. On ADSM servers that support it, a device class that is used to categorize tape devices that support tape cartridges, such as the 3495 Tape Library Dataserver.

cartridge system tape (CST). The base tape cartridge media used with 3480 or 3490 Magnetic Tape Subsystems. When specified as a media type in ADSM, CST identifies standard length tape. Contrast with *enhanced capacity cartridge system tape*.

central scheduler. A function that allows an administrator to schedule client operations and administrative commands. The operations can be scheduled to occur periodically or on an explicit date. See *client schedule* and *administrative command schedule*.

CID. Configuration Installation and Distribution.

client. A program running on a PC, workstation, file server, LAN server, or mainframe that requests services of another program, called the server. There are three types of ADSM clients: administrative, backup-archive, and space management. See *administrative client*, *backup-archive client*, and *space management client*.

Client Access/400. A software product that supports advanced program-to-program communications (APPC) in the DOS, OS/2, and Microsoft Windows environments and provides a set of end user services.

client domain. The set of drives, file systems, or volumes selected by a backup-archive client user during a backup or archive operation.

client migration. The process of copying a file from a client node to ADSM storage and replacing the file with

a stub file on the client node. The process is controlled by the user and by space management attributes in the management class. See also *space management*.

client node.   A file server or workstation on which the backup-archive client program has been installed, which has been registered to the server.

client node definition.   Server control information that includes the client's user ID, password, contact information, policy domain, file compression status, deletion authority, and whether the user ID is locked from the server. A client node definition can be exported from a source server so that it can be imported to a target server at a later date.

client node session.   A period of time in which a user can communicate with a server to perform backup, archive, restore, or retrieval requests. Contrast with *administrative session*.

client options file.   A file that a client can edit, containing a default set of processing options that identify the server, communication method, backup and archive options, space management options, and scheduling options. Also called the *dsm.opt* file.

client polling scheduling mode.   A client/server communication technique where the client queries the server for work.

client schedule.   A database record that describes the planned processing of a client operation during a specific time period. The client operation can be a backup, archive, restore, or retrieve operation, a client operating system command, or a macro. See also *administrative command schedule*.

client/server.   A system architecture in which one or more programs (clients) request computing or data services from another program (server).

client system options file.   A file, used on UNIX clients, containing a default set of processing options that identify the ADSM servers to be contacted for services. This file also specifies communication methods and options for backup, archive, space management, and scheduling. Also called the *dsm.sys* file. See also *client user options file*.

client user options file.   A user-created file, used on UNIX clients, containing a default set of processing options that identify the server, communication method, backup and archive options, space management options,

and scheduling options. Also called the *dsm.opt* file. See also *client system options file*.

closed registration.   A registration process in which an administrator must register workstations as client nodes with the server. Contrast with *open registration*.

collocation.   A process that attempts to keep all data belonging to a single client node on a minimal number of sequential access media volumes within a storage pool. The purpose of collocation is to minimize the number of volumes that must be accessed when a large amount of data must be restored.

command line interface.   A type of user interface where commands are specified on the command line when the backup-archive or administrative client is started. Contrast with *graphical user interface*.

commit.   To make changes permanent in the database files. Changes made to the database files are not permanent until they are committed.

Common Programming Interface Communications (CPI-C).   A programming interface that allows program-to-program communication using SNA LU6.2. See also *Systems Network Architecture*.

Common User Access (CUA).   Guidelines for the dialog between a human and a workstation or terminal. One of the three SAA architectural areas.

communication manager.   A component of OS/2 that allows a workstation to connect to a host computer and use the host resources as well as the resources of other personal computers to which the workstation is attached, either directly or through a host.

communication method.   The method used by a client and server for exchanging information.

communication protocol.   A set of defined interfaces that allow computers to communicate with each other.

compression.   The process of saving storage space by eliminating empty fields or unnecessary data to shorten the length of the file. In ADSM, compression can occur at a workstation before files are backed up or archived to server storage. On some types of tape drives, hardware compression can be used.

Configuration Installation and Distribution (CID).   IBM's term for capabilities to automate installation. CID-enabled products are capable of unattended, remote installation.

contextual help. A type of online help that provides specific information for each selectable object, menu choice, notebook tab, field, and control or push button in a window.

conversion. On VM servers, the process of changing from WDSF/VM to ADSM.

copy group. A policy object that contains attributes that control the generation, destination, and expiration of backup and archive files. There are two kinds of copy groups: backup and archive. Copy groups belong to management classes. See also *frequency*, *destination*, *mode*, *serialization*, *retention*, and *version*.

copy status. The status of volume copies defined to the database or recovery log. The copy status can be synchronized, stale, off-line, or undefined.

copy storage pool. A named set of volumes that contains copies of files that reside in primary storage pools. Copy storage pools are used to back up the data stored in primary storage pools. A copy storage pool cannot be a destination for a backup copy group, an archive copy group, or a management class (for space-managed files). See *primary storage pool* and *destination*.

CPI-C. Common Programming Interface Communications.

CST. Cartridge system tape.

CUA. Common User Access.

D

daemon. In the AIX operating system, a program that runs unattended to perform a standard service. Some daemons are triggered automatically to perform their tasks; others operate periodically.

daemon process. In the AIX operating system, a process begun by the root user or by the root shell that can be stopped only by the root user. Daemon processes generally provide services that must be available at all times, such as sending data to a printer.

damaged file. A file for which ADSM has detected data-integrity errors.

DASD. Direct access storage device.

database. A collection of information about all objects managed by the server, including policy management objects, users and administrators, and client nodes.

database audit. A utility that checks for and optionally corrects inconsistent database references.

database backup series. One full backup of the database, plus up to 32 incremental backups made since that full backup. Each full backup that is run starts a new database backup series. A backup series is identified with a number.

database backup trigger. A set of criteria that defines when and how database backups are run automatically. The criteria determine how often the backup is run, whether the backup is a full or incremental backup, and where the backup is stored.

database buffer pool. Storage that is used as a cache to allow database pages to remain in memory for long periods of time, so that the server can make continuous updates to pages without requiring input or output (I/O) operations from external storage.

database dump. A utility that copies database entries to media for later reload in case a catastrophic error should occur.

database load. A utility that copies database entries from media to a newly installed database.

database volume. A volume that has been assigned to the database.

dataserver. See *Tape Library Dataserver*.

data set. See *linear data set*.

data storage. The primary and copy storage pools used by the server to store users' files: backup versions, archive copies, and files migrated from client nodes (space-managed files). Synonymous with *server storage*. See *primary storage pool*, *copy storage pool*, *storage pool volume*, and *volume*.

DDM. Distributed Data Management.

default management class. A management class assigned to a policy set, which is used to govern backed up or archived files when a user does not explicitly bind a file to a specific management class.

definition. Server control information that includes administrator, client node, and policy definitions. A definition can be exported from a source server to

external media so that it can be imported to a target server at a later date.

deletion exit.   An installation-wide exit that informs a tape management system or operator that the server has deleted a sequential access media volume from its database.

delimiter.   (1) A character used to indicate the beginning and end of a character string. (2) A character that groups or separates words or values in a line of input.

density.   On MVS and VM servers, a device class attribute that identifies the bits per inch that can be stored on tape reels. ADSM supports 1600 and 6250 bits per inch (bpi).

desktop.   On-screen representation of a desk top.

desktop client.   The group of clients supported by ADSM that are not UNIX-based and are not OpenEdition MVS. For example, a DOS client is a desktop client.

destination.   A copy group or management class attribute that specifies the primary storage pool to which a client file will be backed up, archived, or migrated. At installation, ADSM provides storage destinations named BACKUPPOOL, ARCHIVEPOOL, and SPACEMGPOOL.

device class.   A named group of storage devices. Each device class has a unique name and represents a device type of disk, file, optical disk, or tape.

device configuration file.   A file that contains information about defined device classes, and, on AIX servers, defined libraries and drives. The file can be created by using an ADSM command or by using an option in the server options file. The information is a copy of the device configuration information in the ADSM database.

device driver.   A collection of subroutines that control the interface between I/O device adapters and the processor.

device type.   A category of storage device. Each device class must be categorized with one of the supported device types, for example, DISK or CARTRIDGE.

direct access storage device (DASD).   A device in which access time is effectively independent of the location of the data.

DISK.   A device class that is defined by ADSM at installation. It is used to categorize disk drives, such as 3390 DASD or 3380 DASD.

diskette.   A small, magnetic disk enclosed in a jacket.

disk operating system (DOS).   An operating system used in IBM PC, PS/2, and compatible computers.

Distributed Data Management (DDM).   A feature of the System Support Program Product that allows an application program (client) to use server program functions to work on files that reside in a remote system.

DLL.   Dynamic link library.

DLT.   Digital linear tape.

domain.   See *policy domain* or *client domain*.

DOS.   Disk operating system.

drive.   A device used to read and write data on a medium such as a disk, diskette, or tape.

dsm.opt file.   See *client options file* and *client user options file*.

dsmserv.opt.   See *server options file*.

dsm.sys file.   See *client system options file*.

dynamic.   A copy group serialization value that specifies that ADSM accepts the first attempt to back up or archive a file regardless of whether the file is modified during the backup or archive process. See also *serialization*. Contrast with *shared dynamic*, *shared static*, and *static*.

dynamic link library.   A file containing executable code and data bound to a program at load time or run time, rather than during linking. The code and data in a dynamic link library can be shared by several applications simultaneously.

E

ECCST.   Enhanced capacity cartridge system tape.

enhanced capacity cartridge system tape (ECCST). Cartridge system tape with increased capacity that can only be used with 3490E tape subsystems. Contrast with *cartridge system tape*.

error log.   A character file written on random access media that contains information about errors detected by the server or client.

estimated capacity.   The available space, in megabytes, of a storage pool.

Ethernet.   A data link protocol and LAN that interconnects personal computers and workstations via coaxial cable.

event.   Administrative commands or client operations that are scheduled to be executed at a particular time.

event record.   A database record that describes actual status and results for events.

exclude.   The process of identifying files or directories in an include-exclude list to prevent these objects from being backed up whenever a user or schedule issues an incremental or selective backup operation, or to prevent these objects from being migrated off the client node via ADSM space management.

exclude-include list.   See *include-exclude list*.

exit.   To execute an instruction within a portion of a computer program in order to terminate the execution of that portion.

exit machine.   On a VM server, a virtual machine that runs the mount and deletion installation-wide exits on VM systems.

expiration.   The process by which files are identified for deletion because their expiration date or retention period has passed. Backed up or archived files are marked expired by ADSM based on the criteria defined in the backup or archive copy group.

expiration date.   On MVS, VM, and VSE servers, a device class attribute used to notify tape management systems of the date when ADSM no longer needs a tape volume. The date is placed in the tape label so that the tape management system does not overwrite the information on the tape volume before the expiration date.

export.   The process of copying administrator definitions, client node definitions, policy definitions, server control information or file data to external media.

export/import facility.   See *import/export facility*.

extend.   The process of increasing the portion of available space that can be used to store database or recovery log information. Contrast with *reduce*.

F

file data.   File space definitions, authorization rules, backed up files, archive copies, and space-managed files. File data can be exported from a source server to external media so that it can be imported to a target server at a later date.

file record extent.   The extent of the file enumerated in number of records.

file space.   A logical space in a client's storage that can contain a group of files. For clients on systems such as OS/2, a file space is a logical partition and is identified by a volume label. For clients on systems such as AIX and UNIX, a file space can consist of any subset of directories and subdirectories stemming from a virtual mount point. Clients can restore, retrieve, or delete their file spaces from ADSM server storage. ADSM does not necessarily store all the files from a single file space together, but can identify all the files in server storage that came from a single file space.

File Transfer Protocol (FTP).   In TCP/IP, the protocol that makes it possible to transfer data among hosts and to use foreign hosts indirectly.

format.   A device class attribute that specifies the recording format used to read or write to sequential access media, for example to cartridge tape.

frequency.   A copy group attribute that specifies the minimum interval, in days, between incremental backups.

FTP.   File Transfer Protocol.

full backup.   An ADSM function that copies the entire database. A full backup begins a new database backup series. Contrast with *incremental backup*. See *database backup series*.

fuzzy copy.   A backup version or archive copy of a file that might not accurately reflect what is currently in the file because ADSM backed up or archived the file while the file was being modified.

G

general help.   A type of online help that provides an overview of the function of the window.

graphical user interface (GUI).   A type of user interface that takes advantage of a high-resolution monitor, including some combination of graphics, the object-action paradigm, the use of pointing devices, menu bars, overlapping windows, and icons. See *windowed interface*. Contrast with *command line interface*.

group of mirrored volumes.   One, two, or three volume copies defined to the database or recovery log. Each volume copy in the group contains exactly the same portion of the database or recovery log. See *mirroring*.

GUI.   Graphical user interface.

H

handle.   A data structure that is a temporary local identifier for an object. A handle identifies an object at a specific location by binding it.

HDA.   Head-disk assembly.

head-disk assembly (HDA).   A field replaceable unit in a direct access storage device containing the disks and actuators.

help index.   A type of online help that provides an alphabetic listing of all help topics.

hierarchical storage management (HSM) client.   A program that runs on workstations to allow users to maintain free space on their workstations by migrating and recalling files to and from ADSM storage. The HSM client allows use of ADSM space management functions. Synonymous with *space management client*.

high migration threshold.   A percentage of the storage pool capacity that identifies when ADSM can start migrating files to the next available storage pool in the hierarchy. Contrast with *low migration threshold*. See *server migration*.

HP-UX.   Hewlett-Packard UNIX operating system. HP-UX is one of the operating systems that ADSM supports as a client environment and a server environment.

HSM client.   Hierarchical storage management client.

I

import.   The process of copying administrator definitions, client node definitions, policy definitions, server control information or file data from external media to a target server.

import/export facility.   The facility that allows system administrators to copy definitions and file data from a source server to external media to move or copy information between servers. Any subset of information can be imported to a target server from the external media.

inactive version.   A backup version of a file for which a more recently backed up version exists. Inactive backup versions are eligible for expiration processing according to the management class assigned to the file. Contrast with *active version*.

include-exclude file.   On UNIX clients, a file containing statements that ADSM uses to determine whether to back up or migrate certain files, and to determine the associated management classes to use for backup, archive, and space management. See *include-exclude list*.

include-exclude list.   A group of include and exclude option statements in a file. ADSM uses the statements to determine whether to back up or migrate certain files, and to determine the associated management classes to use for backup, archive, and space management. The exclude options identify files that should not be backed up or migrated off the client node. The include options identify files that are exempt from the exclusion rules, or assign a management class to a file or group of files for backup, archive, or space management services. The include-exclude list is defined either in the include-exclude file (for UNIX clients) or in the client options file (for other clients).

inconsistencies.   Any discrepancy between the information recorded in the database about backed up or archived files and the actual data associated with backed up or archived files residing in server storage.

incremental backup.   (1) A function that allows users to back up files or directories that are new or have changed since the last incremental backup. With this function, users can back up files or directories from a client domain that are not excluded in the include-exclude list and that meet the requirements for

frequency, mode, and serialization as defined in the backup copy group of the management class assigned to the files. Contrast with *selective backup*. (2) An ADSM function that copies only the pages in the database that are new or changed since the last full or incremental backup. Contrast with *full backup*. See *database backup series*.

internal mounting facility. On a VM server, a VM facility that allows the server to request tape mounts by sending a message to a mount operator. The message is repeated until the tape is mounted or until the mount wait time is exceeded.

inter-user communication vehicle (IUCV) facility. On a VM server, a VM communication method used to pass data between virtual machines and VM components.

IPX/SPX. Internetwork Packet Exchange/Sequenced Packet Exchange. IPX/SPX is Novell NetWare's communication protocol.

IUCV. Inter-user communication vehicle.

K

KB. Kilobyte.

kernel. The part of an operating system that performs basic functions such as allocating hardware resources.

kernel extension. A program that modifies parts of the kernel that can be customized to provide additional services and calls. See *kernel*.

kilobyte (KB). 1024 bytes.

L

LAN. Local area network.

length. A device class attribute that specifies the length of cartridge tape by specifying one of the following media types: CST for standard length tape or ECCST for double length tape.

library. (1) A repository for demountable recorded media, such as magnetic tapes. (2) In ADSM, a collection of one or more drives, and possibly robotic devices (depending on the library type), which can be used to access storage volumes. (3) In the AS/400 system, a system object that serves as a directory to

other objects. A library groups related objects, and allows the user to find objects by name.

linear data set. A type of MVS data set that ADSM uses for the database, the recovery log, and storage pools. The data set must be preallocated using VSAM IDCAMS and formatted by ADSM for its use. See *minidisk*.

load. See *mount*.

local area network (LAN). A network in which a set of devices are connected to one another for communication and that can be connected to a larger network.

log pool size. The size of an area in memory used to store recovery log pages.

logical volume. The combined space from all volumes defined to either the database or the recovery log. In ADSM, the database is one logical volume and the recovery log is one logical volume.

low migration threshold. A percentage of the storage pool capacity that specifies when ADSM can stop the migration of files to the next storage pool. Contrast with *high migration threshold*. See *server migration*.

M

macro file. An optional file that contains one or more administrative commands and is invoked from an administrative client.

management class. A policy object that users can bind to each file to specify how the server manages the file. The management class can contain a backup copy group, an archive copy group, and space management attributes. The copy groups determine how the ADSM server manages backup versions or archive copies of files. The space management attributes determine whether files are eligible for migration from client nodes to ADSM storage, and under what conditions. See also *copy group*, *binding* and *rebinding*.

mask. A pattern of characters that controls the keeping, deleting, or testing of positions of another pattern of characters or bits.

maximum extension. Specifies the maximum amount of storage space, in megabytes, that you can extend the database or recovery log.

maximum reduction. Specifies the maximum amount of storage space, in megabytes, that you can reduce the database or recovery log.

maximum utilization. The highest percentage of assigned capacity used by the database or recovery log.

MB. Megabyte.

megabyte (MB). (1) For processor storage and real and virtual memory, 2^{20} or 1 048 576 bytes. (2) For disk storage capacity and transmission rates, 1 000 000 bytes.

migrate. (1) To move data from one storage pool to the storage pool specified as the next pool in the hierarchy. The process is controlled by the high and low migration thresholds for the first storage pool. See *high migration threshold* and *low migration threshold*. (2) To copy a file from a client node to ADSM storage. ADSM replaces the file with a stub file on the client node. The process is controlled by the include-exclude list and by space management attributes in management classes.

migration. The process of moving data from one storage location to another. See *client migration* and *server migration*.

minidisk. A logical subdivision of a VM physical disk that provides storage on contiguous cylinders of DASD. On a VM server, a minidisk can be defined as a disk volume that can be used by the database, recovery log, or a storage pool. See also *linear data set*.

mirroring. A feature that protects against data loss within the database or recovery log by writing the same data to multiple disks at the same time. Mirroring supports up to three exact copies of each database or recovery log volume. See *group of mirrored volumes*.

mm. Millimeter.

mode. A copy group attribute that specifies whether to back up a file that has not been modified since the last time the file was backed up. See *modified* and *absolute*.

modified. A backup copy group mode value indicating that a file is considered for incremental backup only if it has changed since the last backup. A file is considered changed if the date, size, owner, or permissions have changed. See *mode*. Contrast with *absolute*.

Motif. A graphical user interface that performs window management and contains a high level toolkit for application program development. It provides an icon view of the UNIX file system. Also known as X-Windows/Motif or Motif X—Toolkit.

mount. To place a data medium (such as a tape cartridge) on a drive in a position to operate.

mount exit. On a VM server, an installation-wide exit (DSMMOUNT EXEC) that requests tape mounts on behalf of the server on VM systems.

mount limit. A device class attribute specifying the maximum number of volumes that can be simultaneously accessed from the same device class, that is, the maximum number of mount points. See *mount point*.

mount operator. On a VM server, a VM user ID that can receive tape mount messages from the server.

mount point. A logical drive through which ADSM accesses volumes in a sequential access device class. For a device class with a removable media device type (for example, CARTRIDGE), a mount point is a logical drive associated with a physical drive. For a device class with the device type of FILE, a mount point is a logical drive associated with an I/O stream. The number of mount points for a device class is determined by the mount limit for that class. See *mount limit*.

mount request. A server request to mount a sequential access media volume so that data can be read from or written to the sequential access media.

mount retention period. A device class attribute that specifies the maximum amount of time, in minutes, that the server retains a mounted sequential access media volume that is not being used before it dismounts the sequential access media volume.

mount wait period. A device class attribute that specifies the maximum amount of time, in minutes, that the server waits for a sequential access volume mount request to be satisfied before canceling the request.

Multiple Virtual Storage (MVS). One of the family of IBM operating systems for the System/370 or System/390 processor, such as MVS/ESA. MVS is one of the supported server environments.

MVS. Multiple Virtual Storage.

N

Named Pipes. A communication protocol that is built into the OS/2 operating system. It can be used to establish communications between an ADSM/2 server and OS/2 clients. The client and ADSM/2 server must reside on the same system.

NetBIOS. Network Basic Input/Output System.

network adapter. A physical device, and its associated software, that enables a processor or controller to be connected to a network.

Network Basic Input/Output System (NetBIOS). An operating system interface for application programs used on IBM personal computers that are attached to the IBM Token-Ring Network.

Network File System (NFS). A protocol defined by Sun Microsystems that extends TCP/IP network file services. NFS permits remote node files to appear as though they are stored on a local workstation.

Networking Services/DOS (NS/DOS). A software product that supports advanced program-to-program communications (APPC) in the DOS and Microsoft Windows 3.1 environments. With NS/DOS, communications applications on your workstation "talk to" partner applications on other systems that support APPC.

NFS. Network File System.

node. A unique name used to identify a workstation to the server. See also *client node.*

notebook. A graphical representation that resembles a spiral-bound notebook that contains pages separated into sections by tabbed divider-pages. A user can "turn" the pages of a notebook to move from one section to another.

notify operator. A VM user ID that specifies an operator who receives messages about severe errors and abnormal conditions.

O

object. A collection of data managed as a single entity.

offsite volume. A removable media volume that is at a location where it cannot be mounted for use.

OpenEdition MVS. MVS/ESA services that support an environment within which operating systems, servers, distributed systems, and workstations share common interfaces. OpenEdition MVS supports standard application development across multivendor systems and is required to create and use applications that conform to the POSIX standard.

open registration. A registration process in which users can register their own workstations as client nodes with the server. Contrast with *closed registration*.

Operating System/2 (OS/2). An operating system used in IBM PC AT, PS/2, and compatible computers. OS/2 is one of the supported client environments and one of the supported server environments.

operator privilege class. An administrative privilege class that allows an administrator to issue commands that control the operation of the server. This privilege class allows disabling or halting the server to perform maintenance, enabling the server, canceling server processes, and managing tape.

optical disk. A disk that contains data readable by optical techniques.

optical drive. A drive mechanism that rotates an optical disc.

optical library. A disk storage device that houses optical disk drives and optical disks, and contains a mechanism for moving optical disks between a storage area and optical disk drives.

OS/2. Operating System/2.

OS/400. Operating System/400.

owner. The owner of backup-archive files sent from a multiuser client node, such as AIX.

P

page. (1) A block of instructions, data, or both. (2) In ADSM, a unit of space allocation within database volumes. (3) In a virtual storage system, a fixed block that has a virtual address and is transferred as a unit between real storage and auxiliary storage.

paging. (1) The action of transferring instructions, data, or both, between real storage and external page storage. (2) Moving data between memory and a mass storage device as the data is needed.

pattern-matching expression. A string expression that uses wildcard characters to specify one or more ADSM objects. See also *wildcard character*.

PC Support/400. A software product that supports advanced program-to-program communications (APPC) in the DOS, OS/2, and Microsoft Windows environments and provides a set of end user services.

platform. The operating system environment in which a program runs.

policy definition. Server control information that includes information about policy domains, policy sets (including the ACTIVE policy set), management classes (including the default management class), copy groups, schedules, and associations between client nodes and schedules. A policy definition can be exported from a source server so that it can be imported to a target server at a later date.

policy domain. A policy object that contains policy sets, management classes, and copy groups that is used by a group of client nodes. See *policy set*, *management class*, and *copy group*.

policy privilege class. An administrative privilege class that allows an administrator to manage policy objects, register client nodes, and schedule client operations (such as backup services) for client nodes. Administrators can be authorized with unrestricted or restricted policy privilege. See *unrestricted policy privilege* or *restricted policy privilege*.

policy set. A policy object that contains a group of management class definitions that exist for a policy domain. At any one time there can be many policy sets within a policy domain but only one policy set can be active. See *management class* and *active policy set*.

premigration. For an HSM client, the process of copying files that are eligible for migration to ADSM storage, but leaving the original file intact on the local system.

primary storage pool. A named set of volumes that ADSM uses to store backup versions of files, archive copies of files, and files migrated from client nodes via ADSM space management. A primary storage pool may be backed up to a copy storage pool either automatically or by command. See *destination* and *copy storage pool*.

privilege class. A level of authority granted to an ADSM administrator. ADSM has five privilege classes: system, policy, storage, operator, and analyst. The privilege class determines which ADSM administrative tasks the administrator can perform. For example, an administrator with system privilege class can perform any administrative task.

programmable workstation communication services (PWSCS). A product that provides transparent high performance communications between programs running on workstations or on host systems.

protection status. A device class attribute that specifies whether to update the RACF profile to identify which users have access to cartridge tapes associated with this device class on MVS servers.

PWSCS. Programmable workstation communication services.

Q

QIC. Quarter-inch cartridge (a type of magnetic tape media).

R

random access media. Any volume accessed in a nonsequential manner. In ADSM, volumes are accessed in a nonsequential manner if they reside in the DISK device class.

randomization. The percentage of the startup window that the server can use to randomize start times for different client nodes associated with a schedule.

rebinding. The process of associating a file with a new management class name. For example, rebinding occurs when the management class associated with a file is deleted. See *binding*.

recall. A function that allows users to access files that have been migrated from their workstations to ADSM storage via ADSM space management. Contrast with *migrate*.

reclamation. A process of consolidating the remaining data from many sequential access media onto a single new sequential access media.

reclamation threshold. A value that specifies a percentage of space on sequential access media volumes that can be occupied by reclaimable space. The remainder of the space is for active data. (Space becomes reclaimable when files are expired.)

recovery log. A log of updates that are about to be written to the database. The log can be used to recover from system and media failures.

recovery log buffer pool. Used to hold new transactions records until they can be written to the recovery log.

reduce. The process of freeing up enough space to allow you to delete a volume from the database or recovery log. Contrast with *extend*.

REEL. On ADSM servers that support it, a device class that is used to categorize tape devices that support tape reels, such as the 3420 9-track tape device.

register. Defines a client node or administrator who can access the server. See *registration*.

registration. The process of identifying a client node or administrator to the server.

reply operator. On a VM server, a VM user ID that specifies an operator who will reply to tape mount requests by the server.

restore. The process of returning a backup copy to an active storage location for use. ADSM has processes for restoring its database, storage pools, storage pool volumes, and users' backed-up files. For example, users can copy a backup version of a file from the storage pool to the workstation. The backup version in the storage pool is not affected. Contrast with *backup*.

restricted policy privilege. An administrative privilege class that enables an administrator to manage policy objects only for the policy domains for which the administrator has been authorized.

restricted storage privilege. An administrative privilege class that enables an administrator to control the allocation and use of storage resources only for the storage pools for which the administrator has been authorized.

retention. The amount of time, in days, that inactive backed up or archived files will be retained in the storage pool before they are deleted. The following copy group attributes define retention: retain extra versions, retain only version, retain version.

retention period. On an MVS server, a device class attribute that specifies how long files are retained on sequential access media. When used, ADSM passes this information to the MVS operating system to ensure that other tape management systems do not overwrite tape volumes that contain retained data.

retrieve. A function that allows users to copy an archive copy from the storage pool to the workstation. The archive copy in the storage pool is not affected. Contrast with *archive*.

RLIO. Record Level Input/Output.

rollback. To remove changes that were made to database files since the last commit point.

root. In the AIX and UNIX environments, the user name for the system user with the most authority.

root user. In the AIX and UNIX environments, an expert user who can log in and execute restricted commands, shut down the system, and edit or delete protected files. Also called the *superuser*.

routing choice. A choice in a pull-down menu that, when selected, brings up another window. See also *action choice*.

S

SAA. Systems Application Architecture.

schedule. A database record that describes scheduled client operations or administrative commands. See *administrative command schedule* and *client schedule*.

scheduling mode. The type of scheduling operation set for the server and client. ADSM supports two scheduling modes for client operations: client-polling and server-prompted.

scratch volume.   A volume that is available for ADSM use. The volume is labeled, is either blank or contains no valid data, and is not defined to ADSM.

SCSI.   Small computer system interface.

selective backup.   A function that allows users to back up specific files or directories from a client domain. With this function, users can back up files or directories that are not excluded in the include-exclude list and that meet the requirement for serialization as defined in the backup copy group of the management class assigned to the files. Contrast with *incremental backup*.

sequential access media.   Any volume that is accessed in a sequential manner, as opposed to a random manner. In ADSM, volumes are accessed sequentially if they reside in a device class other than DISK.

serialization.   A copy group attribute that specifies what ADSM does if files are modified during back up or archive processing. The value of this attribute determines whether processing continues, is retried, or is stopped. See *static*, *dynamic*, *shared static*, and *shared dynamic*.

server.   A program that provides services to other programs (clients).

server migration.   The process of moving data from one storage pool to the next storage pool as controlled by the high and low migration thresholds. See *high migration threshold* and *low migration threshold*.

server options file.   A file that specifies processing options for communication methods, tape handling, pool sizes, language, and date, time, and number formats.

server program.   The program that provides backup, archive, space management, and administrative services to clients. The server program must be at the necessary level to provide all of these services.

server-prompted scheduling mode.   A client/server communication technique where the server contacts the client when work needs to be done.

server storage.   The primary and copy storage pools used by the server to store users' files: backup versions, archive copies, and files migrated from client nodes (space-managed files). Synonymous with *data storage*. See *primary storage pool*, *copy storage pool*, *storage pool volume*, and *volume*.

session resource usage.   The amount of wait time, CPU time, and space used or retrieved during a client session.

shared dynamic.   A copy group serialization value that specifies that a file must not be modified during a backup or archive operation. ADSM attempts to retry the backup or archive operation a number of times; if the file is in use during each attempt, ADSM will back up or archive the file on its last try even though the file is in use. See also *serialization*. Contrast with *dynamic*, *shared static*, and *static*.

shared static.   A copy group serialization value that specifies that the file must not be modified during backup or archive. ADSM will retry the backup or archive operation a number of times; if the file is in use during each attempt, ADSM will not back up or archive the file. See also *serialization*. Contrast with *dynamic*, *shared dynamic*, and *static*.

shell.   In the AIX and UNIX environments, a software interface between a user and the operating system of a computer. Shell programs interpret commands and user interactions on devices such as keyboards, pointing devices, and touch-sensitive screens and communicate them to the operating system.

signal.   (1) A simple method of communication between two processes. One process can inform the other process when an event occurs. (2) In operating system operations, a method of inter-process communication that simulates software interrupts.

signal handler.   A subroutine called when a signal occurs.

SMIT.   System Management Interface Tool.

SNA LU6.2.   Systems Network Architecture Logical Unit 6.2.

socket.   (1) An endpoint for communication between processes or applications. (2) A pair consisting of TCP port and IP address, or UDP port and IP address.

space-managed file.   A file that is migrated from and recalled to a client node via ADSM space management.

space management.   The process of keeping sufficient free storage space available on a client node by migrating files to ADSM storage. The files are migrated based on criteria defined in management classes to which files are bound, and the include-exclude list.

Synonymous with *hierarchical storage management*. See also *migration*.

space management client.   A program that runs on workstations to allow users to maintain free space on their workstations by migrating and recalling files to and from ADSM storage. Synonymous with *hierarchical storage management client*.

SPACEMGPOOL.   A disk storage pool defined by ADSM at installation. It can be the destination for files that are migrated from client nodes via ADSM space management. See *storage pool*.

stale copy status.   Specifies that a volume copy is not available to the database or recovery log.

STANDARD copy group.   A backup or archive copy group that is defined by ADSM at installation. See *copy group*.

STANDARD management class.   A management class that is defined by ADSM at installation. See *management class*.

STANDARD policy domain.   A policy domain that is defined by ADSM at installation. See *policy domain*.

STANDARD policy set.   A policy set that is defined by ADSM at installation. See *policy set*.

stanza.   A group of lines in a file that together have a common function or define a part of the system. Stanzas are usually separated by blank lines or colons, and each stanza has a name.

startup window.   A time period during which a schedule must be initiated.

static.   A copy group serialization value that specifies that the file must not be modified during backup or archive. If the file is modified during the attempt, ADSM will not back up or archive the file. See also *serialization*. Contrast with *dynamic*, *shared dynamic*, and *shared static*.

storage hierarchy.   A logical ordering of primary storage pools, as defined by an administrator with system privilege. Generally, the ordering is based on the speed and capacity of the devices that the storage pools use. In ADSM, the storage hierarchy is defined by

identifying the *next* storage pool in a storage pool definition. See *storage pool*.

storage management services.   A component that allows a central system to act as a file backup and archive server for local area network file servers and workstations.

storage pool.   A named set of storage volumes that ADSM uses to store client data. A storage pool is either a primary storage pool or a copy storage pool. See *primary storage pool* and *copy storage pool*.

storage pool volume.   A volume that has been assigned to an ADSM storage pool. See *volume*, *copy storage pool*, and *primary storage pool*.

storage privilege class.   An administrative privilege class that allows an administrator to control the allocation and use of storage resources for the server, such as monitoring the database, recovery log, and server storage. Administrators can be authorized with unrestricted or restricted storage privilege. See *restricted storage privilege* or *unrestricted storage privilege*.

stub file.   A file that replaces the original file on a client node when the file is migrated from the client node to ADSM storage.

superuser.   See *root user*.

synchronized copy status.   Specifies that the volume is the only volume copy or is synchronized with other volume copies in the database or recovery log. When synchronized, mirroring has started.

system privilege class.   An administrative privilege class that allows an administrator to issue all server commands.

Systems Application Architecture (SAA).   Software interfaces, conventions, and protocols that provide a framework for designing and developing applications that are consistent across systems.

Systems Network Architecture (SNA).   A set of rules for data to be transmitted in a network. Application programs communicate with each other using a layer of SNA called advanced program-to-program communications (APPC).

T

tape. A recording medium consisting of a long, narrow, flexible strip with a magnetic coating wound onto a reel or into a cartridge. See *cartridge* and *tape reel*.

tape library. (1) A term used to refer to a collection of tape cartridges. (2) An automated device that performs tape cartridge mounts and demounts without operator intervention.

Tape Library Dataserver. An automated tape library consisting of mechanical components, cartridge storage frames, IBM tape subsystems, and controlling hardware and software. The tape library dataserver performs tape cartridge mounts and demounts without operator intervention.

tape reel. A cylinder with flanges on which magnetic tape is wound. Devices such as the 3420 9-track tape device support tape reels. Contrast with *cartridge*.

tape volume prefix. A device class attribute that is the high-level-qualifier of the file name or the data set name in the standard tape label.

task help. A type of online help that provides a list of tasks that can be completed with a selected object. When you select a task, the help provides step-by-step information on how to complete the task.

TCP/IP. Transmission Control Protocol/Internet Protocol.

Telnet. In TCP/IP, the protocol that opens the connection to the system.

Transmission Control Protocol/Internet Protocol (TCP/IP). A set of communication protocols that support peer-to-peer connectivity functions for both local and wide area networks.

trusted communication agent. A program that performs communication tasks on behalf of the client or server, and ensures the security of the communications.

U

unit name. On an MVS server, a device class attribute that specifies a group of tape devices used with the MVS server. A unit name can be a generic device type, an esoteric unit name, or a physical device.

unrestricted policy privilege. An administrative privilege class that enables an administrator to manage policy objects for any policy domain.

unrestricted storage privilege. An administrative privilege class that enables an administrator to control the database, recovery log, and all storage pools.

utilization. The percent of assigned capacity used by the database or recovery log at a specific point of time.

V

validate. The process of ensuring that the active policy set contains a default management class and reporting on copy group definition errors.

version. The maximum number of backup copies retained for files and directories. The following copy group attributes define version criteria: versions data exists and versions data deleted.

Virtual Machine (VM). One of the family of IBM operating systems for the System/370 or System/390 processor, including VM/ESA, VM/XA, VM/SP, and VM/HPO. VM is one of the supported server environments.

Virtual Storage Extended (VSE). One of the family of IBM operating systems for the System/370 or System/390 processor, including VSE/ESA. VSE is one of the supported server environments.

VM. Virtual Machine.

volume. The basic unit of storage for the database, recovery log, or a storage pool. A volume can be an LVM logical volume, a standard file system file, a tape cartridge, or an optical cartridge. Each volume is identified by a unique volume identifier. See *database volume*, *scratch volume*, and *storage pool volume*.

volume history file. A file that contains information about: volumes used for database backups and database dumps; volumes used for export of administrator, node, policy, or server data; and sequential access storage pool volumes that have been added, reused, or deleted. The information is a copy of the same types of volume information in the ADSM database.

volume set. An entire image of the database or recovery log, as displayed on the administrative graphical user interface.

VSE.   Virtual Storage Extended.

W

WDSF/VM.   Workstation Data Save Facility/Virtual Machine.

wildcard character.   A character or set of characters used to specify an unknown number or set of characters in a search string.   Also called *pattern-matching character*.

window.   A part of a display screen with visible boundaries in which information is displayed.

windowed interface.   A type of user interface that is either a graphical user interface or a text based interface.   The text based interface maintains a close affinity to the graphical user interface, including action bars and their associated pull-downs and windows.   See *graphical user interface*.

workstation.   A personal computer system capable of maintaining data files.

Workstation Data Save Facility/Virtual Machine (WDSF/VM).   The predecessor product to ADSTAR Distributed Storage Manager.

WORM.   A type of optical media that can only be written to and cannot be erased.

X

X Windows.   A network transparent windowing system developed by MIT.   It is the basis for other products, such as Enhanced X Windows which runs on the AIX operating system.

Index

A

ACTIVATE POLICYSET command   28
activating
   policy set   28
   server   9
activating server options   41
add a library to the system   15, 16
administrative client
   description   xii
   ending   12
   installing   1, 46
   name and password   12
   starting   12
administrative command-line interface   37
administrative graphical user interface (GUI)   37
administrator
   granting authority to   24
   registering   24
ADSM
   installing   1
   setting up   9
ADSM Server Utilities   37
   adding tape drives   13, 15, 16
   creating disk storage pool volumes   11
   description   37
   getting started   9
   initializing a tape volume   18
   start server as service   9
advanced tasks   x
application program interface   52
archive
   description   viii, xii, 9
   to tape   14, 20, 25
archive copy group   viii, 26
archive pool size   11
associating
   client with a policy domain   28
   granting   24
   schedule with clients   30
automated library
   description   13
automating
   administrative commands   33
   client operations   xviii, 30

B

backup
   description   viii, xii, xiv, 9
   selective   10
   to tape   14, 20, 25
backup copy group   viii, 26
backup pool size   11
backup-archive client
   description   xii
   install   1
   register   25

C

cfgserv.log   5
client communications   47
client configuration icon   38
client installation   1, 46
client node
   description   9
   registration   25
client options, modifying   46
client scheduler, starting   32
client, starting the Windows administrative   12
collocation   14
command-line interface, administrative
   quitting   40
   starting   40
   using   40
command, DSMSERV INSTALL   7
commands, administrative
   ACTIVATE POLICYSET   28
   DEFINE ASSOCIATION   30
   DEFINE DEVCLASS   15, 17
   DEFINE DRIVE   15, 16
   DEFINE LIBRARY   15, 16
   DEFINE SCHEDULE   30, 33
   DEFINE STGPOOL   16, 17
   DEFINE VOLUME   19
   EXPIRE INVENTORY   29, 33
   QUERY EVENT   33, 34
   QUIT   40
   REGISTER NODE   25
   UPDATE COPYGROUP   26
   UPDATE NODE   28

Communicating Your Comments to IBM

ADSTAR Distributed Storage Manager
for Windows NT**
Quick Start
Version 2

Publication No. GC35-0235-00

If you especially like or dislike anything about this book, please use one of the methods listed below to send your comments to IBM. Whichever method you choose, make sure you send your name, address, and telephone number if you would like a reply.

Feel free to comment on specific errors or omissions, accuracy, organization, subject matter, or completeness of this book. However, the comments you send should pertain to only the information in this manual and the way in which the information is presented. To request additional publications, or to ask questions or make comments about the functions of IBM products or systems, you should talk to your IBM representative or to your IBM authorized remarketer.

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

If you are mailing a readers' comment form (RCF) from a country other than the United States, you can give the RCF to the local IBM branch office or IBM representative for postage-paid mailing.

- If you prefer to send comments by mail, use the RCF at the back of this book.

- If you prefer to send comments by FAX, use this number:

    - United States and Canada: 520 799-2906
    - Other countries: (1) 520 799-2906

    The contact department is 61C/031.

- If you prefer to send comments by electronic mail, use one of the following addresses:

    - Internet: starpubs@vnet.ibm.com (or starpubs at vnet.ibm.com)
    - IBMLink from U.S.A.: STARPUBS at SJEVM5
    - IBMLink from Canada: STARPUBS at TORIBM
    - IBM Mail Exchange: USIB3VVD at IBMMAIL

Make sure to include the following in your note:

- Title and publication number of this book
- Page number or topic to which your comment applies.

Readers' Comments — We'd Like to Hear from You

ADSTAR Distributed Storage Manager
for Windows NT**
Quick Start
Version 2

Publication No. GC35-0235-00

Overall, how satisfied are you with the information in this book?

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	□	□	□	□	□

How satisfied are you that the information in this book is:

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	□	□	□	□	□
Complete	□	□	□	□	□
Easy to find	□	□	□	□	□
Easy to understand	□	□	□	□	□
Well organized	□	□	□	□	□
Applicable to your tasks	□	□	□	□	□

Please tell us how we can improve this book:

Thank you for your responses. May we contact you? □ Yes □ No

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

_____   _____
Name   Address

_____   _____
Company or Organization

_____   _____
Phone No.

Readers' Comments — We'd Like to Hear from You
GC35-0235-00

IBM®

Cut or Fold
Along Line

Fold and Tape   **Please do not staple**   Fold and Tape
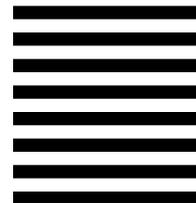
NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST-CLASS MAIL   PERMIT NO. 40   ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

International Business Machines Corporation
Information Development
Department 61C
9000 South Rita Road
TUCSON AZ 85775-4401

Fold and Tape   **Please do not staple**   Fold and Tape

Cut or Fold
Along Line

GC35-0235-00

IBM®

Program Number: 84H09189

Spine information:

IBM **ADSTAR Distributed Storage Manager**
for Windows NT**   Quick Start   *Version 2*