

ADSTAR Distributed Storage Manager
for Sun** Solaris**



Quick Start

Version 2

Note!

Before using this information and the product it supports, be sure to read the general information under Appendix E, "Notices" on page 71.

This book is also available in a softcopy form that can be viewed with the Adobe Acrobat reader licensed program.

First Edition (March 1997)

This edition applies to Version 2 Release 6 of ADSM for Sun Solaris, program number 5639-B22, and to any subsequent release until otherwise indicated in new editions or technical newsletters. Make sure you are using the correct edition for the level of the product.

Order publications through your IBM representative or the IBM branch office serving your locality. Publications are not stocked at the address below.

A form for readers' comments is provided at the back of this publication. If the form has been removed, address your comments to:

IBM Corporation
Information Development, Department 61C
9000 South Rita Road
Tucson, AZ 85744-0001, U.S.A.

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© **Copyright International Business Machines Corporation 1995, 1997. All rights reserved.**

Note to U.S. Government Users — Documentation related to restricted rights — Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract with IBM Corp.

Contents

Before You Begin	vii
What You Should Know before Reading This Publication	vii
Conventions Used in This Book	vii
Types of Information	vii
Where Tasks Are Performed	viii
How to Use This Book	viii
Installation and Setup Results	ix
Basic Tasks	xi
Advanced Tasks	xii
Before You Continue	xiii
ADSM Concepts	xiv
How ADSM Can Change the Way You Work	xiv
System-Managed Storage	xiv
Storage Devices and Media	xviii
Automating Client Operations	xxi
Chapter 1. Installing ADSTAR Distributed Storage Manager	1
Upgrading from ADSM Version 1 to ADSM Version 2	2
Returning from ADSM Version 2 to ADSM Version 1	3
Installing the ADSM Server Program (First Time Installation)	5
Using the ADSM Server Utilities Window	7
Viewing the Online Books with Adobe Acrobat Reader	8
Installing the ADSM Clients for Sun Solaris	10
Verifying Your Installation	13
Removing ADSM from Your System	15
Removing ADSM for Sun Solaris Server	15
Removing ADSM for Sun Solaris Clients	15
Chapter 2. Setting Up the Server	17
Starting the Server	17
Starting the Sun Solaris Administrative Client	17
Changing the Administrative Client Password	18
Stopping the Server	19
Ending an Administrative Client Session	19
Creating Additional Database and Recovery Log Volumes	20
Creating Additional Disk Storage Pool Volumes	21
Registering Your ADSM Licenses	22
Assigning Additional Administrators	23
Registering Additional Backup-Archive Client Nodes	24
Chapter 3. Customizing Your ADSM System	25
Using Tape and Optical Devices with ADSM	25
Attaching the Devices	27
Adding a Manual Tape Library	35
Adding an Automatic Tape Library	36

Displaying Information About Your Devices	38
Preparing Tape Volumes	38
Including Tape in the Server Storage Hierarchy	40
Verifying the Tape Device Environment	41
Creating a New Policy	42
Invoking Expiration Processing	46
Creating a Client Include-Exclude List	46
Automating Client Operations	47
Automating Administrative Commands	50
Protecting ADSM and Client Data	52
Appendix A. Using ADSM Administrative Interfaces	53
Using the Graphical User Interface	53
Defining an Object	53
Changing an Object	53
Deleting an Object	54
Using the Administrative Client Command-Line Interface	54
Using Interactive Mode	54
Appendix B. Setting Server and Client Communications	57
Server Communication Options	57
TCP/IP Communication Method	58
Shared Memory Communication Method	59
Client Communication Options	59
TCP/IP Communication Method	59
Shared Memory Communication Method	61
Appendix C. Applying ADSM Updates	63
Internet, IBMLink, and CompuServe Assistance	63
Installing an ADSM Sun Solaris Server PTF	63
Removing an ADSM Sun Solaris PTF	64
Appendix D. Statement of Service and Where to Look for Information	67
Statement of Service	67
Enhanced Program Services: Support Line	68
CompuServe	68
Internet	68
ADSM Publications	68
IBM Redbooks	70
Software Developer's Program	70
Do You Have Comments or Suggestions?	70
Translations	70
Appendix E. Notices	71
Trademarks	71
Glossary	73

Index 89

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 Sun Solaris
- 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:

- Sun Solaris 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.

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:



Sun Solaris command line



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

After you complete the tasks in Chapter 1, “Installing ADSTAR Distributed Storage Manager” on page 1, “Installing the ADSM Clients for Sun Solaris” on page 10, and Chapter 2, “Setting Up the Server” on page 17, your ADSM system will look like this:

Standard 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.

Database and Recovery Log

A 5MB database volume (db.dsm) and a 9MB recovery log volume (log.dsm)

Disk Storage Pools

Three disk storage pools and two storage pool volumes:

Storage Pool Name	Volume Name and Size
BACKUPPOOL	sbkup00.dsm (10MB)
ARCHIVEPOOL	sarch00.dsm (5MB)
SPACEMGPOOL	none

Clients

Installed and registered administrative and backup-archive clients.

The ADSM installation includes one Sun Solaris administrative client and one Sun Solaris backup-archive client with the following defaults:

Administrative Client

Name *admin*
Password *admin*

Backup-Archive Client

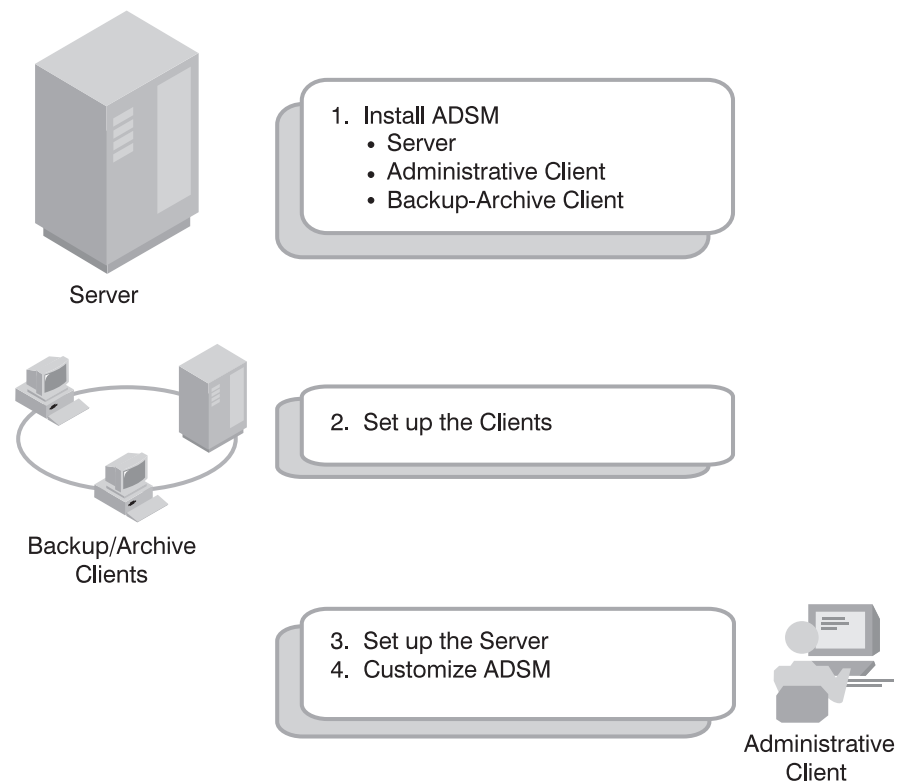
Name *client*
Password *client*

Chapter 3, “Customizing Your ADSM System” on page 25 describes how to do following:

- 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
- 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

Basic Tasks

The following figure 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: installing the server, the Sun Solaris administrative client, and the Sun Solaris backup-archive client (page 1)
Information on installing additional clients is in *ADSM Installing the Clients*
- 2** Setting up the Sun Solaris clients (page 10)
- 3** Setting up the server (page 17)
 - a** Starting the server (page 17)
 - b** Starting an administrative client session (page 17)

- c** Stopping the server (page 19)
- d** Stopping an administrative client session (page 19)
- e** Creating additional database and recovery log volumes (page 20)
- f** Creating additional disk storage pool volumes (page 21)
- g** Registering ADSM licenses (page 22)
- h** Assigning administrators (page 23)
- i** Registering client nodes (page 24)

4 Customizing ADSM (page 25)

- a** Using tape devices with ADSM (page 25)
- b** Verifying the tape device environment (page 41)
- c** Creating a new policy (page 42)
- d** Automating Client Operations (page 47)
- e** Automating Administrative Commands (page 50)

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 server 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 on page xvi shows an overview of the ADSM process for storing client data.

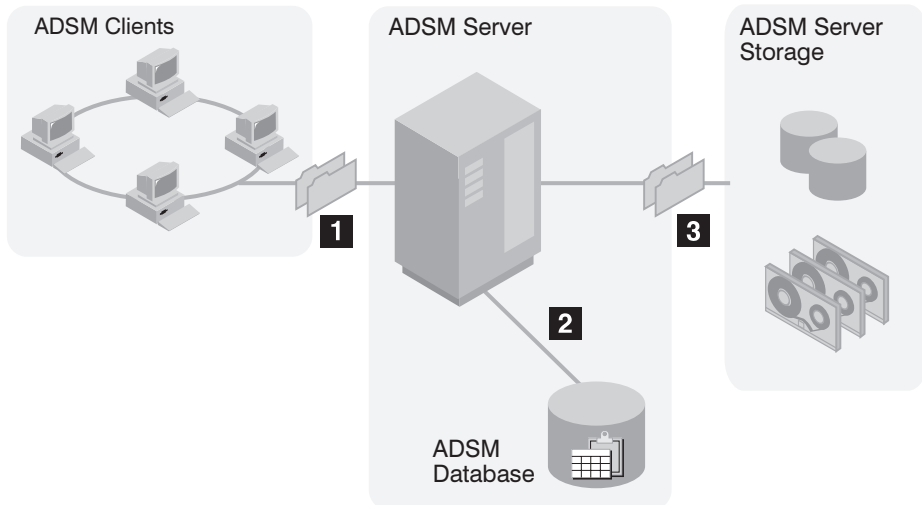


Figure 1. Overview of How ADSM Stores Client Data

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 xvii 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 xx 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.

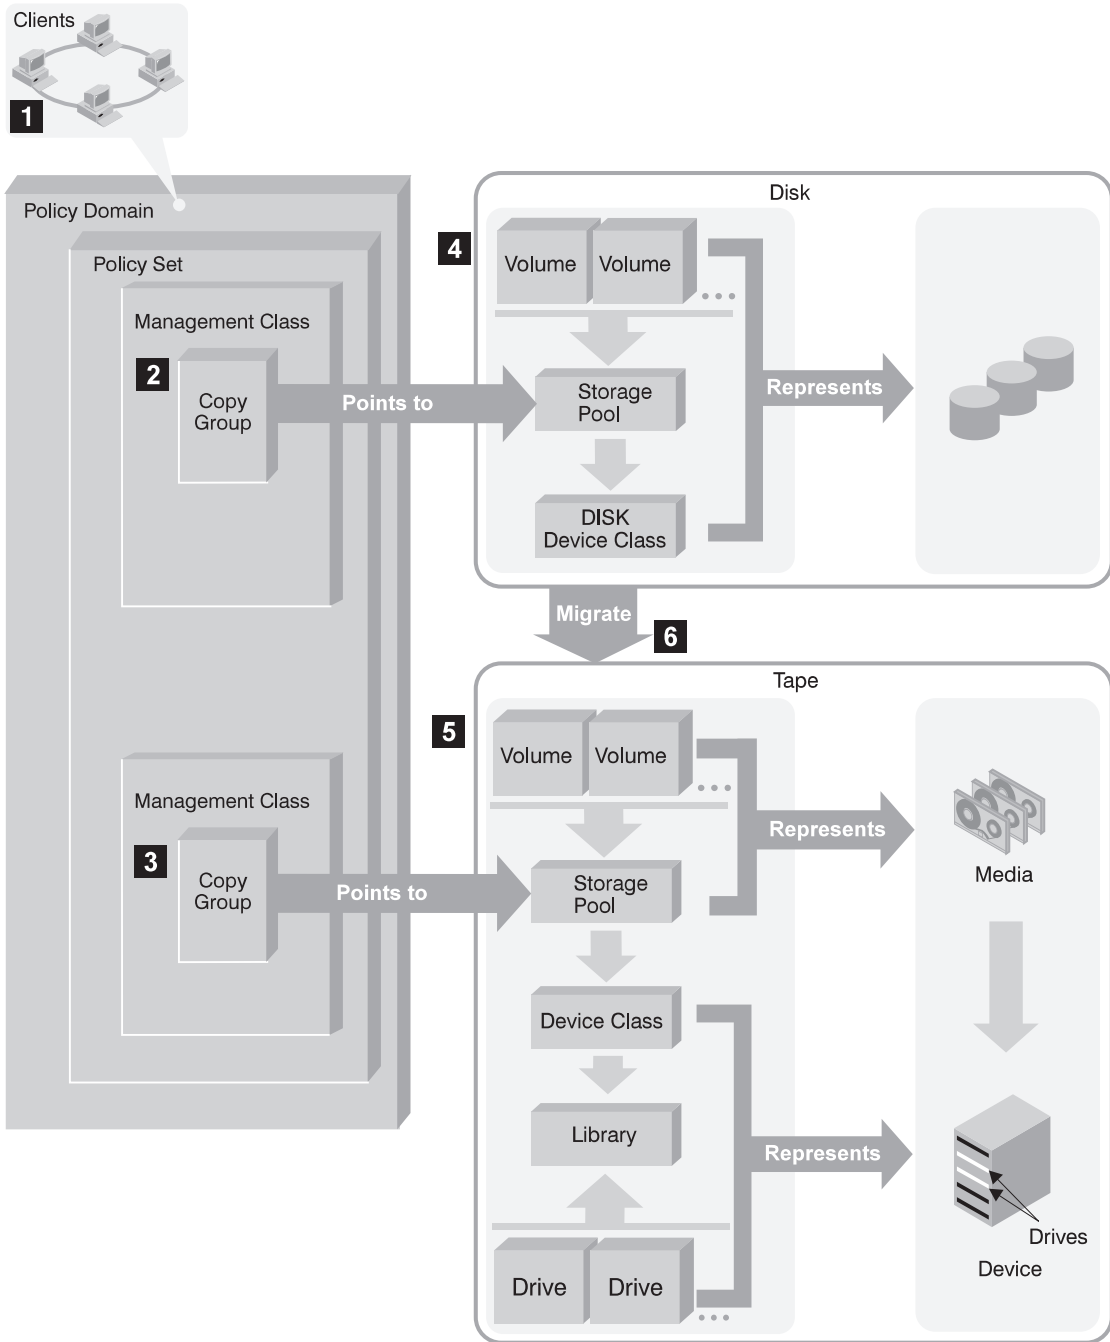


Figure 3. Putting it All Together

Automating Client Operations

You can automate client operations, such as backup and archive. Figure 4 on page xxii 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 xxii.)

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 xxii.) 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 package, the administrative client, and the backup-archive client. The server package includes the ADSM server, an ADSM utility window, and online documentation. This section gives you installation instructions for these components.

When you install the default administrative client and the default backup-archive client, they are registered with the server as shown:

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

The default installation directory for the ADSM Version 2 server package is **/opt/IBMadsm-s/bin**.

The default installation directories for the ADSM Version 2 clients are shown here:

Table 1. Client Components and Directories

Client Component Name	Directory
Backup-archive (command-line and GUI)	/opt/IBMDSMba#
Administrative (command-line)	
Administrative (GUI)	/opt/IBMDSMsa#
API	/opt/IBMDSMa#
Space management (HSM)	/opt/IBMDSMsm#

Note: # is the latest maintenance level (PTF) number.




Before You Install ADSM Version 2

- Ensure that you have an appropriately configured Sun Solaris 2.5.1 or later machine with at least 80MB of free disk storage and 16MB of memory.
 - Ensure that your Sun Solaris system has TCP/IP installed and activated (TCP/IP is the default).
 - Ensure that the ADSM server is halted before installing the Version 2 code.
 - Ensure that you have a CD-ROM device available for installation.
-

Upgrading from ADSM Version 1 to ADSM Version 2

To prevent loss of data from your system if you are migrating from Version 1 to Version 2, follow this procedure:

- 1 Use the online DUMPDB command with your Version 1 server to dump the contents of the server database to sequential media.



```
dsmserv dumpdb
```

This command retrieves database records into a database dump. This process does not access the recovery log. Transactions held in the log database are lost.


Store the output volumes in a safe location. You will need these volumes if you want to return to the Version 1 server after you have installed Version 2.

- 2 Stop the server using the HALT command if it is currently running.



```
halt
```

- 3 Remove the Version 1 server software **IBMads-m-s** by entering:



```
/usr/sbin/pkgrm IBMads-m-s
```

- 4 Install the Version 2 server software. See “Installing the ADSM Server Program (First Time Installation)” on page 5 or enter the following command:



```
/usr/sbin/pkgadd -d /cdrom/ibmadsm2
```

where *-d* indicates device.

This command gives you a list of packages to install. Select IBMads-m-s for the server package.

- 5 Start the server.

When the Version 2 server is started over a database that was written by a Version 1 server program, you must specify the UPGRADEDB parameter with the DSMSERV command.



```
dsmserve upgradedb
```



You need to issue this command only once. Therefore, do not include the UPGRADEDB in any automated programs that start your server.



The server disk file (dsmserve.dsk) is created and identifies the set of recovery log and database volumes that will be used by the server.

- 6 Define ADSM devices as required. See “Using Tape and Optical Devices with ADSM” on page 25 and *Administrator's Guide*.

Returning from ADSM Version 2 to ADSM Version 1



This procedure can result in a loss of client files originally referenced by the Version 1 database. Some client files may have moved or been overwritten or deleted by the Version 2 server program.

If you decide to return to the Version 1 server:

- 1 Use the online BACKUP DB command to take a full backup of your Version 2 database (to be restored if you return to Version 2). Refer to *ADSM Administrator's Guide* for detailed information on this command. For example, to back up the database to FILE, enter:



```
backup db devclass=file type=full
```

- 2 Save a copy of the volume history and device configuration files. For example, to save the volume history in a file called VOLHIST and the device configuration in a file called DEVICES, enter:



```
backup volumehistory filenames=volhist  
backup devconfig filenames=devices
```

- 3 Stop the Version 2 server by using the HALT command after the database backup has successfully completed.

```
ADMIN > halt
```

- 4 Remove the Version 2 server software **IBMads-m-s** by entering:

```
Sun > /usr/sbin/pkgrm IBMads-m-s
```

- 5 Install the Version 1 server software with this command:

```
Sun > /usr/sbin/pkgadd -d /cdrom/ibmadsm IBMads-m-s  
where -d indicates device.
```

- 6 Reinitialize the server database and recovery log using the DSMSERV INSTALL utility.

```
Sun > dsmserv install
```

See *ADSM Administrator's Reference* for details on the DSMSERV command.

- 7 Reload the Version 1 database using the DSMSERV LOADDB utility.

```
Sun > dsmserv loaddb
```

Refer to the *ADSM Administrator's Guide* for details on this command.

- 8 Start the Version 1 server from the **/opt/IBMads-m-s/bin** directory by entering:

```
Sun > ./dsmserv
```

- 9 Issue an AUDIT VOLUME FIX=YES command for each sequential storage pool volume that has a volume type of STGDELETE or STGREUSE to audit all disk storage pool volumes.


- 10 Define the ADSM devices. Refer to *ADSM Administrator's Guide*.

Installing the ADSM Server Program (First Time Installation)

Use this procedure to install the ADSM server program on Sun Solaris.

- 1 Log in as the root user.
- 2 Insert the CD into a CD-ROM drive.
Read the **readme.1st** file in the CD-ROM root directory **/cdrom** before you install ADSM.
- 3 You can make multiple selections at this step, including the Adobe Acrobat reader and the ADSM for Sun Solaris clients. You will need the Adobe Acrobat reader to view the online books. See "Installing the ADSM Clients for Sun Solaris" on page 10 to determine if you want to install the clients at this time.

Enter this command for a list of packages to install:



```
/usr/sbin/pkgadd -d /cdrom/ibmadsm2
```

For the ADSM for Sun Solaris server program, select **IBMads-m-s**. This selection includes the ADSM server, an ADSM utilities window, and the ADSM online documentation.

The **pkgadd** command loads the selected files.

The example here assumes that the CD-ROM name is **/cdrom**, although it may be different in your environment.

Note: Loading the server files can take several minutes.

- 4 During the processing of the above command, the system informs you that:
This package contains scripts which will be executed with super-user (root) permission during the process of installing the package.
Do you want to continue with the installation of this package?


Enter:



y

The installation continues. The ADSM server automatically initializes the database and recovery log and issues a variety of messages indicating its progress.

- 5** When the installation has completed successfully and you are finished with your installation, enter the following commands from a root user command prompt to eject the CD:


```
 cd /  
eject cdrom
```

Note: The ADSM server program files are installed in the **/opt/IBMadsm-s/bin** directory.

- 6** Change to the **/opt/IBMadsm-s/bin** directory and create the server options file from the sample options file.



If this is a reinstall, *do not* recopy the `dsmserv.opt.smp` sample file to the `dsmserv.opt` options file, or you will overwrite any modified options in your existing copy of `dsmserv.opt`.

```
 cd /opt/IBMadsm-s/bin  
cp dsmserv.opt.smp dsmserv.opt
```

Note: The server installation program:

- Installs the ADSM server files
- Installs the online documentation (see “Viewing the Online Books with Adobe Acrobat Reader” on page 8)
- Accesses an ADSM utility window (see “Using the ADSM Server Utilities Window” on page 7)
- Creates a 5MB **db.dsm** database volume
- Creates a 9MB **log.dsm** recovery log volume
- Creates a 10MB **sbkup00.dsm** backup storage pool volume
- Creates a 5MB **sarch00.dsm** archive storage pool volume

See “Creating Additional Database and Recovery Log Volumes” on page 20 and “Creating Additional Disk Storage Pool Volumes” on page 21 for additional information on setting up these volumes.

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

Using the ADSM Server Utilities Window

After you have installed the ADSM server program, you can access the ADSM Server Utilities Window by typing:



The ADSM Server Utilities window accesses configuration and administration utilities for the ADSM administrator. Here is a brief description of the icons on this window.

Getting Started

Offers customer support information, Version 2 highlights, the *ADSM: General Information* manual, help for configuring devices in an autolibrary, and suggestions for things to do after installation.

ADSM Server

Accesses the ADSM console. The console in this interface includes a scrollable server message history and the capability to retrieve and edit commands issued at the command line.

ADSM Administrative Client

Accesses the ADSM administrative interface. Use the administrative interface to communicate with the server to:

- Define ADSM policy
- Define and manage ADSM storage
- Register and manage ADSM clients
- Define and manage ADSM database and recovery log
- Schedule ADSM automatic operations

ADSM Backup-Archive Client

Accesses the backup-archive interface to:

- Back up and restore files
- Archive and retrieve files
- View backup policies

ADSM Client Scheduler

Allows you to start or stop the automatic scheduler daemon that starts the automatic scheduler.

ADSM HSM Client

Accesses the HSM (hierarchical storage management) graphical user interface if you have chosen to install the HSM client.

From this window, you can also access the online documentation by clicking on **Help** and then on **View Books**.

Viewing the Online Books with Adobe Acrobat Reader

You must install the Adobe Acrobat reader from the installation CD-ROM to view the ADSM publications if you do not already have the reader installed on your system. This section offers instructions on installing the Adobe Acrobat reader and a listing of the online books.

Installing the Adobe Acrobat Reader

The Adobe Acrobat files are located in the `/cdrom/ibmadsm2/Acrobat3` directory.

To install Adobe Acrobat, follow this procedure:

- 1 Change to the `/cdrom/ibmadsm2/Acrobat3` directory.



```
cd /cdrom/ibmadsm2/Acrobat3
```

- 2 Run the INSTALL script.



```
./INSTALL
```

The files are installed by default in the `/opt/Acrobat` directory.

- 3 Start the Adobe Acrobat reader by typing `acrobat`.

Note: You can also access the online documentation by clicking on **Help** and then **View Books** at the ADSM Utilities window.

- 4 At the selection bar, type `/opt/IBMads-m-s/ezadsm` to view and select the books.

- 5 Follow the instructions as they appear on your screen.

ADSM Documentation

ADSM V2: Installing the Clients and *ADSM V2: Quick Start* are shipped hard copy with your product.

The ADSM online publications are in PDF format. Table 2 lists the publications and the associated PDF file names that are included in the server installation program.

Table 2. Book Titles and File Names

Book Titles	File Name
Administrator's Guide	a31eas00.pdf
Administrator's Reference	a31ers00.pdf
Quick Start	a31eis00.pdf
Messages	a31emc06.pdf
AFS/DFS Backup Clients	ansafs00.pdf
Installing the Clients	a31eci03.pdf
Trace Facility Guide	a31ect01.pdf
Using the Apple Macintosh Backup-Archive Client	a31ec101.pdf
Using the Application Programming Interface	a31eca01.pdf
Using the DOS Backup-Archive Client	a31ec201.pdf
Using the Lotus Notes Backup Agent	a31ec101.pdf
Using the Microsoft Windows Backup-Archive Client	a31ec601.pdf
Using the Novell Netware Backup-Archive Client	a31ec302.pdf
Using the UNIX Backup-Archive Client	a31ec501.pdf
Using the UNIX HSM Clients	a31ech01.pdf

Installing the ADSM Clients for Sun Solaris

Refer to *Installing the Clients* for installing the administrative and backup-archive clients on other platforms. The clients for other platforms can be installed directly from the CD to the desired client platform. If no CD-ROM drive is available to the client platform, the files can be copied to a file server and distributed.


Use this procedure to install the ADSM for Sun Solaris administrative client and backup-archive clients. You may have already selected to install the clients at step 3 on page 5.

1 Log in as the root user.

2 Insert the CD into a CD-ROM drive.

Read the **readme.1st** file in the CD-ROM root directory **/cdrom** before you install ADSM.

3 Enter the following command:



```
/usr/sbin/pkgadd -d /cdrom/ibmadsm2
```

This command gives you a list of packages to install. You can make multiple selections at this screen.

For the Sun Solaris backup-archive GUI and command-line clients and the administrative command-line client, select **IBMadsmba#**.

See Table 1 on page 1 for the client components.

The **pkgadd** command loads your client file selections.

The example here assumes that the CD-ROM name is **/cdrom**, although it may be different in your environment.

Note: Loading the client files can take several minutes.

4 While processing the command in the previous step, the system asks you several questions:

Do you want to continue with the installation of this package?

Enter:



y

Do you want to install these setuid/setgid files?

Enter:

```
Sun > y
```

Then the system tells you:

This package contains scripts which will be executed with super-user (root) permission during the process of installing the package.

Do you want to continue with the installation of this package?

Enter:

```
Sun > y
```

- 5** When the installation has completed successfully, enter the following commands from a root user command prompt:

```
Sun > cd /  
eject cdrom
```

Note: The client files are installed in the `/opt/IBMSMba#` directory. See Table 1 on page 1 for the value of #.

- 6** Change to the `/opt/IBMSMba#` directory and create the client system options file from the sample options file.

```
Sun > cp dsm.sys.smp dsm.sys
```

- 7** Update client system options file to identify the server and the TCP/IP communication method.



Sun

Edit the `dsm.sys` client options file to change the `SERVERNAME` and `TCPSEVERADDRESS` options to read as follows:

```
SErvername ADSM  
COMMmethod TCPip  
TCPSeveraddress tcp_address  
TCPport 1500
```

During your edit,

- Delete the asterisk (*) in front of each of the above options.
- Replace `tcp_address` with the actual IP address or IP name of your ADSM server.

The options file is neither case-sensitive nor white space-sensitive.

Save the changes made to the client system options file.

- 8** From the `/opt/IBMDSMba#` directory, create the client user options file from the sample options file.



Sun

```
cp dsm.opt.smp dsm.opt
```

- 9** Update client user options file to identify the server and the client node.



Sun

Edit the `dsm.opt` client options file to change the `SERVERNAME` and add the `NODENAME` option to read as follows:

```
SErvername ADSM  
NODename your_node_name
```

During your edit,

- Delete the asterisk (*) in front of each of the above options.
- Replace `your_node_name` with the backup-archive client's node name.

The options file is neither case-sensitive nor white space-sensitive.

Save the changes made to the client user options file.

The installation is complete. You are now ready to verify the installation and begin performing backup-archive client tasks. Continue at "Verifying Your Installation" on page 13.

The installation program:

- Installs the necessary files for server operation
- Installs the online documentation
- Installs a utilities graphical user interface to assist with ADSM tasks
- Registers one client node named CLIENT with a password of CLIENT
- Registers one administrative node named ADMIN with a password of ADMIN
- Creates a 10MB storage pool, BACKUPPOOL, containing the volume SBKUP00.DSM and a 5MB storage pool, ARCHIVEPOOL, containing the volume SARCH00.DSM
- Creates a 9MB recovery log volume, log.dsm
- Creates a 5MB database volume, db.dsm
- Creates a dsmserv.dsk file that contains database and recovery log information in your **/opt/IBMads-m-s** directory

Verifying Your Installation

- 1** Start the ADSM server session by entering from the **/opt/IBMads-m-s/bin** directory:

```
Sun > ./dsmserv
```

The server must be activated and running to run the administrative and backup-archive clients.

If you installed the client component, you can now verify the installation by doing the following:

- 2** Start the administrative client command line session by entering from the **/opt/IBMDSMba#** directory (see Table 1 on page 1 for the client components).

```
Sun > dsmadm
```

The default passwords for the administrative client created at installation are:

Name *admin*
Password *admin*

You can issue query commands at this time. For query command information, enter:

```
ADMIN > help query
```

To exit the administrative client command line session, enter:

 `quit`

3 Start the backup-archive client with one of these methods:

- Double-click on the backup-archive GUI icon:




- Issue the following command:


 `dsm`

When you start the backup-archive client, you must enter a node name and password. The defaults for the backup-archive client created at ADSM installation are:

Name *client*
Password *client*

4 Back up several files. For example, to back up all the files in the `/dev/cd0` directory that begin with **proj**, do the following:

 `dsmc selective /dev/cd0/proj*.*`

- 
- END USER
1. Double-click on the backup-archive GUI icon
 2. Select the entry or entries in **File Systems for Backup-Archive** list that contain the files you want to back up.



Do not select files to back up that exceed your server capacity. A 10MB backup storage pool volume and a 5MB archive storage pool volume were allocated at installation.

5 Click on **Backup**.


6 You will receive a message indicating if the backup has completed successfully.

Removing ADSM from Your System

This section assists you if you need to remove the ADSM server or clients from your system.

Removing ADSM for Sun Solaris Server

Enter the following command if you need to remove the ADSM server package from your system:

```
 /usr/sbin/pkgrm IBMadsm-s
```

A series of informational messages are displayed that indicate the system is obtaining file information, removing selected files, and finally confirming that removal is complete.

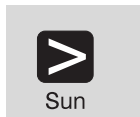


Your database and recovery log files contain information pertaining to backup-archive clients. The database and recovery log are not deleted when you issue the **pkgrm** command.

Use caution if you decide to reinstall the server with the DSMSEV INSTALL command in the same directory. The DSMSEV INSTALL command will destroy any data that previously resided in your database, recovery log, and storage pool volumes.

Removing ADSM for Sun Solaris Clients

Enter the following command if you need to remove the ADSM clients from your system:

```
 /usr/sbin/pkgrm IBMDSMba#  
See Table 1 on page 1 for the value of #.
```

A series of informational messages are displayed that indicate the system is obtaining file information, removing selected files, and finally confirming that removal is complete.

Chapter 2. Setting Up the Server

This section describes how to set up and start your server.

Here are some ADSM terms in this chapter:

archive	To copy a file to a storage pool for long-term storage.
client node	A file server or workstation on which the backup-archive client program has been installed, which has been registered to the server.
storage pool	A named set of volumes for storing client data.

Starting the Server

The ADSM server must be running before any client can connect. There are several ways to start the server:

- Issue the following command from the **/opt/IBMadsm-s/bin** directory:



- From the ADSM GUI, double-click on the **ADSM Server** icon.

You can now issue commands from the server console, or you can work from the Sun Solaris administrative client GUI or command line.

Starting the Sun Solaris Administrative Client

This section describes how to start the Sun Solaris administrative client. 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 administrator client created at ADSM installation are:

Name	<i>admin</i>
Password	<i>admin</i>

To change these values see “Changing the Administrative Client Password” on page 18.

- To start the Sun Solaris **administrative client GUI**, issue the following command:

```
 dsmadm
```

- To start the Sun Solaris **administrative client command line interface** issue the following command:

```
 dsmadm
```




The command line examples in this book are for the Sun Solaris administrative client. Sun Solaris requires continuation characters when a command is on more than one line. Other clients may not accept continuation characters. Check the requirements for your command line client.

You can perform most tasks from either the command line or the GUI (for example, registering client nodes). The description of these tasks includes both command line and GUI; choose whichever one is more convenient for you. You can perform some tasks only from the command line (for example, defining libraries). For information about the administrative client interfaces, see Appendix A, “Using ADSM Administrative Interfaces” on page 53.

Changing the Administrative Client Password

If you wish, you can change the password of the administrative client (ADMIN) at this point. To change the password to TUCSON, for example, do the following:

```
 update admin admin tucson
```



1. Double click on **Administrators**
2. From the list of administrator names, double click on ADMIN
3. Click on the **Password** tab
4. Enter the following:

New password	TUCSON
Reenter new password	TUCSON
5. Click on the **Apply** push button
6. Close the **ADMIN - Properties** window

Stopping 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:



halt



1. Double click on **Server**
2. Click on **Server** from the menu bar
3. Click on **Halt** from the pull down menu

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

Ending an Administrative Client Session

To end an administrative client session do one of the following:



From a command line client in interactive mode:
quit



From a GUI client, one of the following:

- Double click on the system menu symbol

OR

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

Creating Additional Database and Recovery Log Volumes

The installation procedure creates a default 5MB database volume (DB.DSM) and a default 9MB recovery log volume (LOG.DSM). The database size is largely determined by the number of client files being stored on server storage. The database should be at least large enough to hold one day's worth of client backups. 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 increase the size of the database or recovery log, do the following:


- 1** Format an additional 13MB database volume and 5MB recovery log volume.



```
./dsmfmt -m -db /opt/IBMadsm-s/bin/db2.dsm 13
```

```
./dsmfmt -m -log /opt/IBMadsm-s/bin/log2.dsm 5
```

- 2** Define the database volume and recovery log volumes, and extend the database and recovery log by entering:



```
define dbvolume db2.dsm
```

```
define logvolume log2.dsm
```

```
extend db 13
```

```
extend log 5
```



Define the database volume:

1. Expand **Database**
2. Double click on **Database Volumes**
3. Click on **Edit** from the menu bar
4. Click on **Add Volume** from the pull down menu
5. Enter the following:
Mirrored volumes
Volume 1 db2.dsm
6. Select **Show Extend window next** check box
7. Click on the **Add** push button
8. Click on **Extend** push button
9. Close the **Database Volumes** window



Define the recovery log volume:

1. Expand **Recovery Log**
2. Double click on **Recovery Log Volumes**
3. Click on **Edit** from the menu bar
4. Click on **Add Volume** from the pull down menu
5. Enter the following:
Mirrored volumes
Volume 1 log2.dsm
6. Select **Show Extend window next** check box
7. Click on the **Add** push button
8. Click on **Extend** push button
9. Close the **Recovery Log Volumes** window

Creating Additional Disk Storage Pool Volumes

At server installation, three default disk storage pools and two storage pool volumes are created:

Storage Pool Name	Volume Name and Size
BACKUPPOOL	sbkup00.dsm (10MB)
ARCHIVEPOOL	sarch00.dsm (5MB)
SPACEMGPOOL	none

You can create additional storage pool volumes. For example, to create a 25MB volume named sbkup01.dsm in BACKUPPOOL, do the following:

1 Format the volume:



```
./dsmfmt -m -data /opt/IBMadsm-s/bin/sbkup01.dsm 25
```

2 Define the volume:



```
define volume backuppool sbkup01.dsm
```



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	BACKUPPOOL
Volume name	sbkup01.dsm
6. Click on the **Add** push button
7. Close the **Storage Pool Volumes** window

Registering Your ADSM Licenses

When you install ADSM for Sun Solaris, your system is licensed for the base ADSM support, which includes an unlimited number of administrative clients and one Sun Solaris backup-archive client.

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
5client.lic	5 backup-archive clients
10client.lic	10 backup-archive clients
50client.lic	50 backup-archive clients
network.lic	Clients requiring a network communications method, TCP/IP for example
desktop.lic	Desktop clients (OS/2, Windows, Apple, Novell NetWare, and DOS)
unix.lic	Unix clients
oemvs.lic	Open System Environment clients

spaceman.lic	Space management feature
drm.lic	Disaster Recovery Manager (DRM) feature
devmod1.lic	Device Module 1
devmod2.lic	Device Module 2
devmod3.lic	Device Module 3
devmod4.lic	Device Module 4
devmod5.lic	Device Module 5 (use of secondary server attachment)
devm1to2.lic	Upgrade from Device Module 1 to 2
devm2to3.lic	Upgrade from Device Module 2 to 3
devm3to4.lic	Upgrade from Device Module 3 to 4

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.

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 DESKTOP environment, and allow attaching devices in Device Support Module 3, issue the following commands:

```

ADMIN register license file(50client.lic)
ADMIN register license file(desktop.lic)
ADMIN register license file(network.lic)
ADMIN register license file(devmod3.lic)

```

Assigning 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.

```

ADMIN register admin mark missouri
ADMIN grant authority mark classes=system

```



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
6. Click on the **Password** tab, and in the **Password** and **Reenter password** fields, enter MISSOURI
7. Click on the **Add** push button
8. Close the **Administrators** window

Registering 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.



register node mercedes 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
7. Close the **Nodes** window

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 still does not take advantage of all the options available to you as an ADSM administrator. However, it does provide 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 "Starting the Server" on page 17 for the procedure).

Here are some ADSM terms that you will find in this chapter:

collocation	A process by which the server attempts to keep all files belonging to a client on a minimal number of tape volumes.
device class	A group of storage devices of the same type, for example, disk or cartridge.
expiration	The process by which files in server storage are marked for deletion because their retention period has passed.
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.
validation	The process of ensuring that the active policy set contains a default management class and reporting on copy group definition errors.

Using Tape and Optical Devices with ADSM

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

1. Attach the devices to your system, which includes physically attaching the devices, installing the device drivers, and determining the device names
2. Define the following: a library for the drives, the drives, a device class, and a storage pool associated with the device class
3. Include the storage pool in your hierarchy

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

This section describes the following tasks on the pages indicated:

Task	Page
Attach the devices	27
Add a manual tape library	35
Add an automatic tape library	36
Prepare tape volumes for use	38
Verify the use of the devices.	41

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 8mm tape drives
- An 8mm automatic tape library with two 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.



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

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

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*.

Attaching the Devices

For ADSM to use a device, you must physically attach the devices to the server and install the appropriate device driver on your system. The following sections discuss how this is done for drives in a manual library and for the drives and controller in automated libraries.



IBM 3490E, 3494, 3570, and 3590 devices

Install the IBMtape device driver. For an IBM 3494, install the IBM Tape Library Driver for Sun Solaris, which includes the Imcp daemon. See “Installing Device Drivers for IBM 3490, 3570, and 3590 Tape Drives” on page 30, “Installing Device Drivers for IBM 349X Libraries” on page 31, and *IBM SCSI Device Drivers: Installation and User's Guide*.

Non-IBM tape drives and tape autochangers

Install the ADSM device drivers. When you install ADSM, you must choose whether to install the ADSM device drivers for tape and tape autochanger devices. See “Installing Device Drivers for Non-IBM Tape Drives and Libraries” on page 31 for how to install the device drivers.

Optical devices

Install the ADSM device drivers.

The drivers are started when you start Sun Solaris. The devices must be powered on for Sun Solaris recognize them.

Attaching Drives in a Manual Library

Here is an overview of the procedure for attaching a drive in a manual library to the server system:

- 1** Ensure that the SCSI adapter card is installed in your system.
- 2** Determine the SCSI IDs available on the SCSI adapter card to which you are attaching the device. Find one unused SCSI ID for each drive.
- 3** Follow the manufacturer's instructions to set the SCSI ID for the device to the unused SCSI ID that you found. Usually this means setting switches on the back of the device.

Attention: Each device connected in a chain to a single SCSI bus must be set to a unique SCSI ID. If each device does not have a unique SCSI ID, you may have serious system problems.

- 4** Follow the manufacturer's instructions to attach the device to your server system hardware.

Attention:

- a. Power off your system before attaching a device to prevent damage to the hardware.
- b. You must attach a terminator to the last device in the chain of devices connected on one SCSI adapter card. Detailed instructions should be in the documentation that came with your hardware.

- 5** Install the appropriate device driver. See "Attaching the Devices" on page 27.

- 6** Determine the name for each device, which is needed to define the device to ADSM.

For a tape drive, the device name is a special file name with the form `/dev/rmt/xmt`, where `x` is the number assigned when the drive was set up with its device driver.

Attaching an Autochanger and its Drives

Here is an overview of the procedure for attaching an automated library to the server system:

- 1** Ensure that the SCSI adapter card is installed in your system.

- 2** Determine the SCSI IDs available on the SCSI adapter card to which you are attaching the device. Find one unused SCSI ID for each drive, and one for the library or autochanger controller.

Attention: In some automated libraries, the drives and the autochanger share a single SCSI ID, but have different LUNs. For these libraries, only a single SCSI ID is required. Check the documentation for your device.

- 3** Follow the manufacturer's instructions to set the SCSI ID for the drives and library controller to the unused SCSI IDs that you found. Usually this means setting switches on the back of the device.

Attention: Each device connected in a chain to a single SCSI bus must be set to a unique SCSI ID. If each device does not have a unique SCSI ID, you may have serious system problems.

- 4** Follow the manufacturer's instructions to attach the device to your server system hardware.

Attention:

- a. Power off your system before attaching a device to prevent damage to the hardware.
- b. You must attach a terminator to the last device in the chain of devices connected on one SCSI adapter card. Detailed instructions should be in the documentation that came with your hardware.

- 5** Install the appropriate device drivers. See "Attaching the Devices" on page 27.

- 6** Find the device worksheet that applies to your device. See the device configuration worksheets appendix in the *ADSM Administrator's Guide*.

- 7** Determine the name for each device, which is needed to define the device to ADSM, and record the name on the device worksheet.

The device name for each drive has the form `/dev/rmt/xmt`. The device name for the library controller has the form `/dev/rmt/xb`. In both cases, *x* is the number assigned when the drive and library was set up with its device driver.

Keep the Worksheets: The information you record on the worksheets can help you when you need to perform operations such as adding volumes to an autochanger. Keep them for future reference.

Setting the Library Mode: For ADSM to access a SCSI library, the device must be set for the appropriate mode. The mode that ADSM requires is usually called *random* mode; however, terminology may vary from one device to another. Two examples follow:

- Some libraries have front panel menus and displays that can be used for explicit operator requests. However, if the device is set to respond to such requests, it typically will not respond to requests made by ADSM.

- Some libraries can be placed in *sequential* mode, in which volumes are automatically mounted in drives by using a sequential approach. This mode conflicts with how ADSM accesses the device.

Refer to the documentation for your device to determine how to set it to a mode appropriate for ADSM.

Installing Device Drivers for IBM 3490, 3570, and 3590 Tape Drives

For IBM 3490, 3570, and 3590 devices, see *IBM SCSI Device Drivers: Installation and User's Guide* for how to install the device drivers.

After completing the procedure in the manual, a set of special file names are available for each device. To see the special file names, use the command

```
> ls -la /dev/rmt
```

The output shows that the special files are really symbolic links into the /devices tree of the file system.

- If you are installing the device driver for an IBM 3490 or an IBM 3590 without an ACF, you need to select one special file name for the drive. The name typically has the form:

```
/dev/rmt/xst
```

where *x* is a number. The name may have additional characters (n, b, c) at the end to indicate different operating characteristics.

Enter the device name that you select in the Device Name field for the drive on the worksheet.

- If you are installing the device driver for an IBM 3590 Model B11 with an ACF, or an IBM 3570, you need to select two special file names for the device: one for the drive and one for the medium changer.

The name for the drive typically has the form:

```
/dev/rmt/xst
```

where *x* is a number. The name may have additional characters (n, b, c) to indicate different operating characteristics. Enter the name that you select for the drive in the Device Name field for the drive. For example, to use the compression capability of the drive, enter:

```
/dev/rmt/1stc
```

The name for the medium changer has the form:

```
/dev/rmt/xsmc
```

where *x* is a number. The *smc* stands for SCSI medium changer. Enter this name in the Device Name field for the library's medium changer (robotics). For example:

```
/dev/rmt/0smc
```

Installing Device Drivers for IBM 349X Libraries

For an IBM 3494 Tape Library Dataserver, use the information in *IBM SCSI Device Drivers: Installation and User's Guide* to install the driver. The driver is the IBM Tape Library Driver for Sun Solaris, which consists of the *Imcpd* (a daemon) and other components.

To determine the name to use to define the library to ADSM, check what you entered in the `/etc/ibmatl.conf` file for the library's symbolic name. Use this information to complete the Device Name field on your worksheet.

For example, if the symbolic name for the library in the `/etc/ibmatl.conf` file is 3494a, enter that name on the worksheet in the Device Name field for the library.

The drives that are in the library are set up separately. See "Installing Device Drivers for IBM 3490, 3570, and 3590 Tape Drives" on page 30.

Installing Device Drivers for Non-IBM Tape Drives and Libraries

To use the ADSM device drivers, you must first edit configuration files that reside in the `/usr/kernel/drv` directory, and then issue the `add_drv` (add drive) command. The devices should already be connected to the system and powered on.

If you are configuring a tape library, you must set up the device drivers for both the drives in the library and the library itself (that is, the robotics).

Installing the ADSM Device Driver for Tape Drives: To install the ADSM device driver for tape drives, perform the following steps:

- 1 Verify that the device is connected to your system, and is powered on and active.
- 2 Edit the `/usr/kernel/drv/mt.conf` file. Add one stanza (as shown in the example at the top of the file) for each ADSM tape drive in the system that you want the ADSM server to use.
- 3 Log in as root user.
- 4 Add the drive by entering the following command:

```
> /usr/sbin/add_drv -m '* 0666 bin bin' /usr/kernel/drv/mt
```

- 5 Determine the special file names for the tape drives.

The device special files for the tape drives appear in the `/dev/rmt` directory. Their names have the form `/dev/rmt/xmt`, where *x* is a number. To determine the relationship between a physical device and the special file that was created for it, issue the following command:

```
> ls -l /dev/rmt/*mt
```

The output shows that the special file is really a symbolic link into the /devices tree of the file system. The symbolic link destination has the form /devices/.../mt@x,0:mt, where *x* is the SCSI ID of the device. This information lets you know which of the /dev/rmt/xmt special file names to provide to the server when you issue a DEFINE DRIVE command.

- 6 Record the SCSI IDs and special file names on the worksheets. See the device configuration worksheets appendix in the *ADSM Administrator's Guide*.

Installing the ADSM Device Driver for Tape Libraries: To install the ADSM device driver for tape libraries, perform the following steps:

- 1 Edit the /usr/kernel/drv/lb.conf file. Add one stanza (as shown in the example at the top of the file) for each ADSM autochanger device in the system that you want the ADSM server to use.
- 2 Log in as root user.
- 3 Add the library by entering the following command:

```
> /usr/sbin/add_drv -m '* 0666 bin bin' /usr/kernel/drv/lb
```

- 4 Determine the special file names for the tape libraries.

The device special files for libraries appear in the /dev/rmt directory. Their names have the form /dev/rmt/xlb, where *x* is a number. To determine the relationship between a physical device and the special file that was created for it, issue the following command:

```
> ls -l /dev/rmt/*lb
```

The output shows that the special file is really a symbolic link into the /devices tree of the file system. The symbolic link destination has the form /devices/.../lb@x,y:lb, where *x* is the SCSI ID of the device and where *y* is the LUN (logical unit number) of the device. This information lets you know which of the /dev/rmt/xlb special file names to provide to the server when you issue a DEFINE LIBRARY command.

- 5 Record the SCSI IDs and special file names on the worksheets. See the device configuration worksheets appendix in the *ADSM Administrator's Guide*.

Installing Device Drivers for Optical Devices

To use the ADSM device drivers, you must first edit configuration files that reside in the `/usr/kernel/drv` directory, and then issue the `add_drv` (add drive) command. The devices should already be connected to the system and powered on.

If you are configuring an optical library, you must set up the device drivers for both the drives in the library and the library itself (that is, the robotics) for ADSM to use.

Installing the ADSM Device Driver for Optical Drives: To install the ADSM device driver for optical drives, perform the following steps:

- 1 Verify that the device is connected to your system, and is powered on and active.
- 2 Edit the `/usr/kernel/drv/op.conf` file. Add one stanza (as shown in the example at the top of the file) for each ADSM tape drive in the system that you want the ADSM server to use.
- 3 Log in as root user.
- 4 Add the drive by entering the following command:

```
/usr/sbin/add_drv -m '* 0666 bin bin' /usr/kernel/drv/op
```

- 5 Determine the special file names for the drives.

The device special files for the optical drives appear in the `/dev/rmt` directory. Their names have the form `/dev/rmt/xop`, where `x` is a number. To determine the relationship between a physical device and the special file that was created for it, issue the following command:

```
ls -l /dev/rmt/*op
```

The output shows that the special file is really a symbolic link into the `/devices` tree of the file system. The symbolic link destination has the form `/devices/.../op@x,0:op`, where `x` is the SCSI ID of the device. This information lets you know which of the `/dev/rmt/xop` special file names to provide to the server when you issue a `DEFINE DRIVE` command.

- 6 Record the SCSI IDs and special file names on the worksheets. See the device configuration worksheets appendix in the *ADSM Administrator's Guide*.

Installing the ADSM Device Driver for Optical Libraries: To install the ADSM device driver for a library, perform the following steps:

- 1 Edit the `/usr/kernel/drv/lb.conf` file. Add one stanza (as shown in the example at

the top of the file) for each ADSM autochanger device in the system that you want the ADSM server to use.

- 2** Log in as root user.
- 3** Add the library by entering the following command:

```
> /usr/sbin/add_drv -m '* 0666 bin bin' /usr/kernel/drv/lb
```

- 4** Determine the special file names for the tape libraries.

The device special files for libraries appear in the /dev/rmt directory. Their names have the form /dev/rmt/xlb, where *x* is a number. To determine the relationship between a physical device and the special file that was created for it, issue the following command:

```
> ls -l /dev/rmt/*lb
```

The output shows that the special file is really a symbolic link into the /devices tree of the file system. The symbolic link destination has the form /devices/.../lb@x,y:lb, where *x* is the SCSI ID of the device and where *y* is the LUN (logical unit number) of the device. This information lets you know which of the /dev/rmt/xlb special file names to provide to the server when you issue a DEFINE LIBRARY command.

- 5** Record the SCSI IDs and special file names on the worksheets. See the device configuration worksheets appendix in the *ADSM Administrator's Guide*.

Determining Device Names for ADSM

To identify and work with a removable media device, ADSM needs the device's name. For devices other than the IBM 3494, the name is the special file name. For an IBM 3494, the device name is the symbolic name you entered in the configuration file for the library.

During configuration of devices other than the IBM 3494, one or more device special file names are created. The following table shows the form of the names. In this table, *x* denotes a number (0–7).

Special File Name	Description
/dev/rmt/xmt	Used for tape drives that are supported by the ADSM device driver
/dev/rmt/xlb	Used for SCSI-attached libraries that are supported by the ADSM device driver
/dev/rmt/xop	Used for optical drives that are supported by the ADSM device driver
/dev/rmt/xst	Used for IBM 3490, 3570, and 3590 tape drives, which are supported by the IBMtape device driver
/dev/rmt/xsmc	Used for the medium changer (the robotics) of an IBM 3570 or an IBM 3590 B11 with Automatic Cartridge Facility (ACF) feature, which are supported by the IBMtape device driver

The special file is a symbolic link into the /devices tree of the file system. The symbolic link destination has the form:

```
/devices/.../mt@x,0:mt
```

where *x* is the SCSI ID of the device.

The following are examples of special file names:

Table 3. Examples of Device Special File Names

If the logical file name is:	The device special file name is:
3mt	/dev/rmt/3mt
0lb	/dev/rmt/0lb
1op	/dev/rmt/1op
2st	Drive: /dev/rmt/2st Medium changer: /dev/rmt.2smc (3570 or 3590 B11 with ACF)

Adding a Manual Tape Library

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

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 MANUAL8MM. The two drives, named DRIVE01 and DRIVE02, are at SCSI adaptor 0 and SCSI IDs 1 and 2.



```
define library manual8mm libtype=manual
define drive manual8mm drive01 device=/dev/rmt/1mt
define drive manual8mm drive02 device=/dev/rmt/2mt
```

- 2** Classify the devices 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, TAPE8MM_CLASS).



```
define devclass tape8mm_class devtype=8mm -  
format=8500 library>manual8mm mountlimit=2
```

- 3** Define a storage pool named **TAPE8MM_POOL** for the device class.



```
define stgpool tape8mm_pool tape8mm_class -  
maxscratch=20
```



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 TAPE8MM_POOL
5. Click on the **Device class** tab, and enter the following:
Device class TAPE8MM_CLASS (from the drop-down list)
Maximum scratch volumes 20
6. Click on the **Add** push button
7. Close the **Storage Pools** window


Adding an Automatic Tape Library

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

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 AUTO8MMLIB, whose SCSI library mechanism is at SCSI ID 3. The two drives, named DRIVE05 and DRIVE06, are at SCSI IDs 5 and 6. The element numbers for the drives in an EXB-210 are 82 and 83.



```
define library auto8mmlib libtype=scsi -  
  device=/dev/rmt/31b  
define drive auto8mmlib drive05 device=/dev/rmt/5mt -  
  element=82  
define drive auto8mmlib drive06 device=/dev/rmt/6mt -  
  element=83
```



Element address

The element address is a number that indicates the physical location of a drive within an automated library. ADSM needs the element address to connect the physical location of the drive to the drive's SCSI address. When you define a drive, the element address is required if there is more than one drive in an automated library.

- 2** Classify the drives in the library 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, AUTO8MM_CLASS).



```
define devclass auto8mm_class devtype=8mm -  
  format=8500 library=auto8mmlib mountlimit=2
```

- 3** Define a storage pool named AUTO8MM_POOL for the device class.




```
define stgpool auto8mm_pool auto8mm_class -  
  maxscratch=20
```



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 AUTO8MM_POOL
5. Click on the **Device class** tab, and enter the following:
Device class auto8mm_class (from the drop-down list)
Maximum scratch volumes 20
6. Click on the **Add** push button
7. Close the **Storage Pools** window

Displaying Information About Your Devices

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`

Preparing Tape Volumes

All tape volumes must have 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 standard labels, label them. From an Sun Solaris command prompt, go to the ADSM server directory (for example, `/opt/IBMadsm-s/bin`):



1. Enter the following command:
`dsmlabel -drive=/dev/rmt/1mt`
2. When prompted,
 - a. Insert the new volume into drive **/dev/rmt/1mt**
 - b. Type a 1 to 6 character volume name (for example, **DSM001**).
 - c. Press Enter

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



```
define volume tape8mm_pool 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 TAPE8MM_POOL (from the drop-down list)

Volume name DSM001

6. Click on the **Add** push button
7. Close the **Storage Pool Volumes** window

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 standard labels, label them. From an Sun Solaris command prompt, go to the ADSM server directory (for example, /opt/IBMadsm-s/bin):




1. To specify /dev/rmt/5mt, which is at element address 82 in library /dev/rmt/3lb, enter the following command:

```
dsmlabel -drive=/dev/rmt/5mt,82 -  
-library=/dev/rmt/3lb -search -keep
```

The library searches for available tapes and loads them into the drive one at a time.

2. When prompted,
 - a. Type a 1 to 6 character volume name (for example, DSM001).
 - b. Press Enter

- 2** Check the tape volumes into the automated library.



Ensure that the volumes are in their slots, and issue the CHECKIN LIBVOLUME command:

```
checkin libvolume auto8mm1ib status=scratch  
search=yes
```

The library checks in each volume automatically.

3 Define any volumes that are to be used as private volumes.



```
define volume auto8mm_pool 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 AUTO8MM_POOL (from the drop-down list)

Volume name DSM001

6. Click on the **Add** push button
7. Close the **Storage Pool Volumes** window

Including 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 AUTO8MM_POOL and how to change ARCHIVEPOOL so that data migrates to TAPE8MM_POOL.



```
update stgpool backuppool nextstgpool=auto8mm_pool
```

```
update stgpool archivepool nextstgpool=tape8mm_pool
```




1. Double click on **Storage Pools**
2. Double click on **BACKUPPOOL**
3. Click on the **Migration** tab, go to page 2, and for **Next storage pool**, do the following:
 - a. Deselect (by clicking on) the **None** check box
 - b. Click on AUTO8MM_POOL from the **Next pool** drop-down list
4. Click on the **Apply** push button
5. Close the **BACKUPPOOL - Properties** window
6. Double click on **ARCHIVEPOOL**
7. Click on the **Migration** tab, go to page 2, and for **Next storage pool**, do the following:
 - a. Deselect (by clicking on) the **None** check box
 - b. Click on TAPE8MM_POOL from the **Next pool** drop-down list
8. Click on the **Apply** push button
9. Close the **ARCHIVEPOOL - Properties** window

Verifying the Tape Device Environment

Do the following to verify that the tape devices have been integrated into your ADSM system.

1 Migrate data from a disk storage pool.

Migration from the a storage pool requires tape mounts. The mount messages are directed to the console message queue and to any administrative client that has been started with either the mount mode or console mode option.

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.



```
update stgpool backuppool highmig=0 lowmig=0
```



1. Double click on **Storage Pools**
2. Double click on **BACKUPPOOL**
3. Click on the **Migration** tab, go to page 2, and enter the following:

Migration Thresholds Low	0
Migration Thresholds High	0
4. Click on the **Apply** push button
5. Close the **BACKUPPOOL - Properties** window

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



```
update stgpool backuppool highmig=90 lowmig=70
```



1. Double click on **Storage Pools**
2. Double click on **BACKUPPOOL**
3. Click on the **Migration** tab, go to page 2, and enter the following:

Migration Thresholds Low	70
Migration Thresholds High	90
4. Click on the **Apply** push button
5. Close the **BACKUPPOOL - Properties** window



Controlling when migration occurs

Because migration from a disk to a tape storage pool uses resources such as drives and operators, you may want to ensure that migration occurs at a time that is best for your situation. You can control when migration occurs by leaving the high migration threshold set to 100 most of the time, and lowering it when you want migration to occur.

In the preceding step, you lowered both migration thresholds to 0 to verify that migration occurred. You may not want to empty the disk storage pool by setting the low migration threshold to 0 every time you want migration to occur. Normally, you may want to keep the low threshold at 40%, and vary the high threshold from as high as 90% to as low as 50%. See *ADSM Administrator's Guide* for details.

Creating a New Policy

The ADSM default policy specifies how ADSM manages client files. For example, it specifies that ADSM retains up to two backup versions of active client files (files that exist on the client). Two versions may be satisfactory for most clients. However, assume that a group of clients needs the last ten versions to be kept. You can do either of the following:

- Create a new policy domain and assign these clients to that domain (described in this section).
- Create a new management class within the default policy domain. Then the include-exclude lists for all the affected clients must be updated.



Storing Directly to Tape

The default policy directs that client files are stored directly to disk. You have just defined a tape storage pool and updated your disk storage pools so that they migrate the client data to that tape pool. If you wish, you could define policies that store client files directly to tape by specifying in a copy group a tape pool as the destination.

However, 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.



copy domain standard newdomain



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 Using** from the pull-down menu
5. Enter the following:
Policy domain NEWDOMAIN
6. Click on the **Add** push button
7. Close the **Policy Domains** window

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.



```
update copygroup newdomain standard standard standard -  
type=backup verexists=10
```



1. Expand **Policy Domains**
2. Double click on **Backup Copy Groups**
3. Double click on the STANDARD copy group from policy domain NEWDOMAIN and policy set STANDARD (you cannot update a copy group from policy set ACTIVE)
4. Click on the **Copy control** tab, go to page 2, and enter the following:

Number of backup versions to keep if client data exists	10
--	----
5. Click on the **Apply** push button
6. Close the **Backup Copy Group - Properties** window

3 Validate the STANDARD policy set in NEWDOMAIN.



```
validate policysset newdomain standard
```



1. Expand **Policy Domains**
2. Double click on **Policy Sets**
3. Click on the STANDARD policy set in NEWDOMAIN
4. Click on **Selected** 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
8. Close the **Policy Sets** window



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 the NEWDOMAIN policy domain.



```
activate policyset newdomain standard
```



1. Expand **Policy Domains**
2. Double click on **Policy Sets**
3. Click on the STANDARD policy set in NEWDOMAIN
4. Click on **Selected** 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
8. Close the **Policy Sets** window

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

9. Click on **Backup Copy Groups**
10. View copy group STANDARD for policy set ACTIVE in NEWDOMAIN by double clicking on that line
11. Click on the **Copy control** tab, go to page 2, and verify that the number of backup versions to keep if client data exists is set to 10

- 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 **mercedes**, do either of the following:



update node mercedes domain=newdomain



1. Double click on **Nodes**
2. Double click on **MERCEDES** from the list of node names
3. Select NEWDOMAIN in the **Policy domain** field from the drop-down list
4. Click on the **Apply** push button
5. Close the **CLIENT - Properties** window

ADSM will now keep up to ten versions of backed up files for client node **mercedes** because it is assigned to the NEWDOMAIN policy domain.

Invoking Expiration Processing

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 “Automating Administrative Commands” on page 50). If you schedule expiration processing, set the expiration interval server option to 0.



expire inventory

Creating 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 or, on UNIX clients, an include-exclude file. For example, an include-exclude file should 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. Users can include these statements in their client options file (dsm.opt). 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:

```
excl ude d:\eng\spec\*.*
incl ude d:\eng\spec\*.drw
incl ude 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, the user can also override the default management class.

```
excl ude d:\eng\spec\*.*
incl ude d:\eng\spec\*.drw monthly
incl ude 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.

Automating 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)

In this section, you are guided through the steps necessary to automate backup and archive operations for clients by creating schedules. The example assumes that you have registered three client nodes named **bill**, **mark**, and **mercedes**.

- 1** Define a schedule for an incremental backup operation and associate 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. This schedule calls for a schedule window that:

- Begins on the date the schedule is defined (the default) at 11:00 p.m.
- Lasts for 1 hour (the default)
- Is repeated daily (the default)
- Stays in effect indefinitely (the default)



```
define schedule standard daily_incr action=incremental -
starttime=23:00
```

```
define association standard daily_incr bill,mark,mercedes
```




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	DAILY_INCR

6. Associate nodes by doing the following:
 - a. Click on the **Add nodes** push button
 - 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 click on INCREMENTAL from the **Action** drop-down list
8. Click on the **Timing** tab, and specify the **Start time** as 23:00
9. Click on the **Add** push button
10. Close the **Backup/Archive Schedules** window

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)



```
define schedule standard arc action=archive -  
  objects="c:\* d:\*" starttime=23:00 -  
  period=1 perunits=weeks dayofweek=friday
```

```
define association standard arc bill,mark,mercedes
```



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
 - 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
10. Close the **Backup/Archive Schedules** window

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



```
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 DAILY_INCR schedule runs.

- 4** Because the DAILY_INCR schedule is to run daily, you can verify that it 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*.



```
query event standard daily_incr begindate=today-1
```



1. Expand the **Central Scheduler** object
 2. Double click on **Backup/Archive Events**
 3. Check the status of the DAILY_INCR schedule
-



You can limit the query of events to display only schedules that did not run as expected. For example, you can use the following command daily to see which clients did not run the DAILY_INCR schedule the previous day:

```
query event standard daily_incr begindate=today-1 -  
exceptionsonly=yes
```

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

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

Automating 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)



```
define schedule expproc type=administrative -  
cmd='expire inventory' active=yes starttime=02:00
```



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
8. Close the **Administrative Command Schedules** window

- 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*.



```
query event expproc type=administrative -  
begindate=today-1
```



1. Expand the **Central Scheduler** object
2. Click on **Administrative Command Events**
3. Check the status of the EXPPROC schedule

Protecting 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.

In addition, Disaster Recovery Manager (DRM), an optional feature of ADSM, can assist you in many of the tasks associated with protecting and recovering your data. 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:

- Graphical user interface (GUI) (see page 53)
- Administrative command-line interface (see page 54)

Using the Graphical User Interface

You can manage the server from a graphical user interface.

For instructions on starting the Sun Solaris GUI, see "Starting the Sun Solaris Administrative Client" on page 17.

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 and arrow buttons. Click on the tabs and arrow buttons 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 **Selected** 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 and arrow buttons. Click on the tabs and arrow buttons 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 the Administrative Client Command-Line Interface

The administrative client command-line interface 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 Sun Solaris, the administrative client invokes ADSM by using the `DSMADMC` command. This section shows Sun Solaris 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:

```
dsmadm
```

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:

```
dsmadm -id=smith -password=secret
```

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

```
quit
```

Appendix B. Setting Server and Client Communications

This section offers communications information for your system.



For details about other options, such as date and time formats, buffer sizes, and message queues, see *ADSM Administrator's Reference*.

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

ADSM options shown in this chapter appear in mixed case and do not need to be entered as shown. Uppercase characters represent the minimum characters necessary to process the option. For example, the ADSM option TCPPort can be entered as:

TCPP or tcpp

All of the sample or default entries in the options file are entered as comments, as shown by leading asterisks (*).

When editing the options file, follow these guidelines:

- Options can begin in any column of the input line.
- Only one option can be on each line, and the option cannot span more than one line.
- Keywords can be in upper or lower case. The minimum abbreviation is shown in upper case.
- One or more blank spaces are allowed between keywords and values.
- ADSM interprets lines beginning with an asterisk (*) as comment lines.
- File names must be correctly specified, as required by the operating system.

When editing the options file, be sure to remove the asterisk at the beginning of the line to activate that option.

Server Communication Options

The ADSM for Sun Solaris server uses TCP/IP or SHARED MEMORY or both as the communication method to the clients. These options are specified in the **dsmserv.opt** file.



You can also specify a communication method of NONE in your options file. If COMMMethod NONE is included in your set of communication method options, it overrides all other communication options. This option does not allow users to connect to the server, and it is useful for experimenting with policy commands.

TCP/IP Communication Method

Figure 5 displays an example of a TCP/IP communication method setting.

COMMMETHOD	TCPip
TCPPort	1500
TCPWindowsize	0
TCPNODELAY	YES

Figure 5. 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.

If the server uses both TCP/IP and Shared Memory communication methods, the port numbers for each must differ.

The port number for the client and the server must correspond. If a second server is installed, the port number must differ from the first server.

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 window size is specified in kilobytes. You can specify a value from 0 to 64. 0 indicates that the default window size set for Sun Solaris should be used. Values from 1 to 64 indicate that the window size ranges from 1KB to 64KB. This value must be an integer.

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 YES.

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

NO Specifies that TCP/IP should buffer small ADSM messages

Shared Memory Communication Method

Figure 6 displays an example of a Shared Memory communication method setting.

COMMMETHOD	SHAREdmem
SHMPort	1510

Figure 6. Example of Shared Memory Communication Method Options

SHMPort Specifies the TCP/IP port address of a server when using Shared Memory. All shared memory communications start with a TCP/IP connection. The SHMPORT option lets you specify a different TCP/IP port on which to listen for clients. The default value is 1510, but if you are using port 1510 for some other application, you must specify a different SHMPORT for both the client and the server. You can specify a value from 1024 to 32767.

Client Communication Options

Administrative clients can use the TCP/IP communication protocol or the SHARED MEMORY communication protocol (if the server and client reside on the same machine) to communicate with the server.

You are required to specify the communication method, along with other options, in the client system options file **dsm.sys** for each server.

You are then required to specify which server to use for server services in the client user options file **dsm.opt**.

TCP/IP Communication Method

Transmission Control Protocol/Internet Protocol (TCP/IP) is supported for ADSM clients. Figure 7 shows an example of a TCP/IP communication method setting.

SERVERNAME	server_name
COMMMETHOD	TCPip
TCPSEVERADDRESS	tcp_address
TCPPORT	1500
TCPBUFFSIZE	31
TCPWINDOWSIZE	32

Figure 7. Example of TCP/IP Communication Method Options

TCPSEVERADDRESS Specifies the internet server address for communication using TCP/IP.

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. This value must be the same as that being used for TCP/IP communications on the server. 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.

DOS clients should set a small buffer size.

Specify the size of the buffer in kilobytes. The table here shows the minimum, default, recommended, and maximum values for several clients.

Table 4. TCPBuffsize

Client	Minimum	Recommended	Maximum	Default
UNIX	1	32	32	16 ¹ 31 ²
OS/2	1	32	32	16 ¹ 31 ²
Netware	1	32	32	16 ¹ 31 ²
Windows NT/95	1	32	32	16 ¹ 31 ²
DOS	1	32	32	8 ¹ 31 ²
Others	1	32	32	16 ¹ 31 ²

Note:

- ¹ memory-constrained
- ² NOT memory-constrained

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

DOS clients should set a small window size.

Specify the size of the buffer in kilobytes. The table here shows the minimum, default, recommended, and maximum values for several clients.

Table 5. TCPWindowSize

Client	Minimum	Recommended	Maximum	Default
UNIX (most)	1	64	2048	16 ¹ 32 ²
Sun Solaris	1	64	64	16 ¹ 32 ²
OS/2	1	32	2048	16 ¹ 32 ²
Netware	1	32	64	16 ¹ 32 ²
Macintosh	4	32	128	16 ¹ 32 ²
Windows NT/95	1	32	2048	16 ¹ 32 ²
Others	1	24	24	16 ¹ 32 ²

Note:

- ¹ memory-constrained
- ² NOT memory-constrained

Shared Memory Communication Method

Shared Memory communications is supported between Sun Solaris clients and servers on the same machine. To use Shared Memory communications, TCP/IP must be installed on the machine.

Figure 8 shows an example of a Shared Memory communication method setting used in Sun Solaris.

```

COMMETHOD          SHAREdmem
  SHMPort            1510
  
```

Figure 8. Example of Shared Memory Communication Method Setting Used in Sun Solaris Client

SHMPort Specifies the TCP/IP protocol to be used for the initial connection. The default value is 1510. You can specify a value from 1000 to 32767. This value must be the same as that being used for SHARED MEMORY communications on the server. Contact your system administrator for this value.



These restrictions apply for Shared Memory protocol:

- The user executing the client process (for example, dsm, dsmc) must be running as root or as the same user ID as the processor running the server (dmserv). Some work-arounds, as listed below, exist for these restrictions, but none of the work-arounds eliminate the requirement to authenticate to the server. The work-arounds simply allow non-root users to use the Shared Memory protocol. The user still needs to supply a valid password, if one would be required anyway.

The work-arounds are:

- 1 If you use "PASSWORDACCESS GENERATE" in the **dsm.sys** options file, ADSM uses the dsmtca (Trusted Communications Agent) process. This work-around can be used for the backup-archive clients (dsm and dsmc).
- 2 If you use the administrative clients (dsmadm and dsmadm), you do not use the dsmtca process. However, you can change the ownership of the dsmadm and dsmadm programs to root and set the setuid bit on. For example:

```
chown root.system dsmadm
chmod u+s dsmadm
```

This causes dsmadm to execute as root, which gives it authority to connect. You still need to pass authentication checking for the actual administrative session.

- The Shared Memory protocol has a limit of 50 concurrent sessions. If you try to establish more than 50 sessions at any one time, you may receive out-of-memory errors.
-

Appendix C. Applying ADSM Updates

Application of ADSM maintenance updates, also referred to as Program Temporary Fixes (PTFs), bring your ADSM server up to the current maintenance level.

On Sun Solaris, the fixes are installed in the `/usr/adsmid/opt/IBMadsm-f/bin/dsmserv` directory.

Internet, IBMLink, and CompuServe Assistance

Online help is available through Internet, IBMLink/ServiceLink, and CompuServe.

Internet: If your installation has access to the Internet, the latest server and client PTFs (maintenance updates) are available on the anonymous ftp (file transfer protocol) server. The machine name is `index.storsys.ibm.com`.

IBMLink/ServiceLink: If your installation has access to IBMLink/ServiceLink, an interactive online database program, you can:

- Search for an existing authorized program analysis report (APAR) that is similar to your problem.
- Search for an available program temporary fix (PTF) for the existing APAR.
- Order the PTF if it is available.
- Create an Electronic Technical Response (ETR) problem report to get assistance from an ADSM service representative.

CompuServe: If your installation has access to CompuServe, you can obtain code fixes for the ADSM clients and ADSM servers through the IBMSTORAGE forum libraries.

You can also access general information files, helpful hints, Technical Support Bulletins, and ADSM informational APAR text through the IBMSTORAGE forum.

To access this forum, type **GO ADSM** or **GO IBMSTORAGE** at any ! prompt. You can use this forum to exchange information and ideas with IBM representatives and other ADSM users.

Installing an ADSM Sun Solaris Server PTF



Ensure that the ADSM server is NOT RUNNING when service is being applied. Use the HALT command from the ADSM server console or from an ADSM administrative client. Failure to follow this guideline may result in server crashes.

Each PTF is shipped as a compressed PKGADD-installable file. To install a PTF, complete the following procedure:

- 1** Obtain the compressed PKGADD-installable file, placing it in any suitable directory.
- 2** Uncompress the file using the UNCOMPRESS command.
- 3** Use the `/usr/sbin/pkgadd` command to install the PTF package.
Assuming that the file was downloaded as `/tmp/ptf.pkg.Z` and uncompressed to `/tmp/ptf.pkg`, the customer would use the following PKGADD command:



```
/usr/sbin/pkgadd -d /tmp/ptf.pkg IBMadsm-f
```

IBMadsm-f is the package name used for PTF packages for the ADSM server.



The PTF package must be installed using the same base directory specification as was used for the base (GA) version of the ADSM server. For example, if the **-R** option was used for the initial installation of the ADSM server, then the same **-R** option must be used for all PTF packages applied to the server.


- 4** After the PKGADD command completes successfully, you can restart the ADSM server.

Removing an ADSM Sun Solaris PTF

If you decide to *back out* of a PTF after having applied it to a system, use the `/usr/sbin/pkgrm` command.

Usually, the first PTF that is applied to the system will be given a package instance name of **IBMadsm-f**. The second and third PTF packages will then be given instance names of **IBMadsm-f.2** and **IBMadsm-f.3**. This sequence will continue for further PTF packages. Only the most recent PTF can be backed out at any given time, so you must specify the package instance associated with the most recent PTF installation. Unless you use special overrides when installing a PTF, the most recent one will have the largest instance number (the number after the period in the package instance name).

For example, to remove the PTF whose package instance name is **IBMads-m-f.3**, use the following command:

 `/usr/sbin/pkgrm IBMads-m-f.3`

If you want to revert to a very early version of the ADSM server code, you need to sequentially back out of all PTF packages that were installed beyond the desired server level.

As with installation of a PTF, the ADSM server must be halted before any attempt to back out of a PTF. Failure to follow this guideline may result in server crashes.

Appendix D. Statement of Service and Where to Look for Information

This section describes the service that IBM provides for ADSM and tells you where to look for more information about ADSM.

Statement of Service

IBM Program Services (U.S.A. and Puerto Rico)

IBM Program Service is available for ADSM for Sun Solaris Version 2 until the Service Expiration Date, March 31, 1999.

Defect support is provided through IBM Program Service for this program. A defect is defined as a code-related or documentation-related problem.

A program defect can be reported to the IBM Support Center by electronic mail, facsimile, or mail. IBM will respond in the same medium as the inquiry within two business days if notified by facsimile or electronic mail, and within seven business days if notified by U.S.A. mail (or equivalent). Defect information can be submitted to any of the following:

- Internet Email:
aixsupt@aix.boulder.ibm.com
Subject: AIX Program Services
- Fax:
Number: 1-512-823-7634
Subject: AIX Program Services
- U.S.A. Mail Address:
IBM Corporation
Attention: AIX Program Services
Mail Drop 2900
11400 Burnet Road
Austin, TX 78758

IBM Program Service will respond by issuing one of the following:

- Corrective service information, such as corrected documentation
- Code corrections
- Notice of availability of corrected code
- Consideration as enhancement for future release
- Restriction
- Bypass
- Other notice

IBM does not guarantee service results, or represent or warrant that all errors or program defects will be corrected.

To save you time and allow us to better assist you, please include the following information when reporting defects:

- The ADSM server version, release, modification, and service level number. You can get this information by entering `q` status from the ADSM server console.
- ADSM client's version, release, modification, and service level number. You can get this information by entering `dsmc` from the command line or by using the *Product Information* option of the graphical user interface help menu.
- The communication protocol, version, and release number you are using (for example, TCP/IP).
- The activity you were doing when the problem occurred, listing the steps you followed before the problem occurred.
- The exact text of any error messages.

Enhanced Program Services: Support Line

Value added services are available through the AIX Support Family of Services. Customers are encouraged to purchase the Support Family options that best meet their needs. The Support Family options add voice and electronic access with an increased response criteria and optional extended hours of support. For a complete list of the AIX Support Line options, contact the AIX Support Family Marketing Center at:

1-800-CALL-AIX (225-5249)

CompuServe

You can get additional information electronically by accessing the IBMSTORAGE Forum on CompuServe. To access this forum, enter:

GO IBMSTORAGE

at any ! prompt.

Internet

You can get additional information electronically through an anonymous ftp server. The host name for this machine is: `index.storsys.ibm.com`.

There is also an ADSM-L newsgroup, which is on `listserv@marist.edu`. This bulletin board is implemented by a third party, which IBM supports on a best-effort basis only.

ADSM Publications

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

Short Title	Publication Title	Order Number
Hardcopy Publications		
ADSM Messages	<i>ADSTAR Distributed Storage Manager: Messages</i>	SH35-0133
ADSM License Information	<i>ADSTAR Distributed Storage Manager for Sun Solaris: License Information</i>	SC35-0261
ADSM Administrator's Guide	<i>ADSTAR Distributed Storage Manager for Sun Solaris: Administrator's Guide</i>	GC35-0263
ADSM Administrator's Reference	<i>ADSTAR Distributed Storage Manager for Sun Solaris: Administrator's Reference</i>	GC35-0264
ADSM Using the UNIX HSM Clients	<i>ADSTAR Distributed Storage Manager: Using the UNIX Hierarchical Storage Management Clients</i>	SH26-4030
ADSM V2 Using the Apple Macintosh Backup-Archive Client	<i>ADSTAR Distributed Storage Manager Version 2: Using the Apple Macintosh Backup-Archive Client</i>	SH26-4051
ADSM Using the UNIX Backup-Archive Clients	<i>ADSTAR Distributed Storage Manager Version 2: Using the UNIX Backup-Archive Clients</i>	SH26-4052
ADSM V2 Using the OS/2 Backup-Archive Client	<i>ADSTAR Distributed Storage Manager Version 2: Using the OS/2 Backup-Archive Client</i>	SH26-4053
ADSM V2 Using the DOS Backup-Archive Client	<i>ADSTAR Distributed Storage Manager Version 2: Using the DOS Backup-Archive Client</i>	SH26-4054
ADSM V2 Using the Novell NetWare Backup-Archive Client	<i>ADSTAR Distributed Storage Manager Version 2: Using the Novell NetWare Backup-Archive Client</i>	SH26-4055
ADSM V2 Using the Microsoft Windows Backup-Archive Clients	<i>ADSTAR Distributed Storage Manager Version 2: Using the Microsoft Windows Backup-Archive Clients</i>	SH26-4056
ADSM Using the Lotus Notes Backup Agent	<i>ADSTAR Distributed Storage Manager: Using the Lotus Notes Backup Agent</i>	SH26-4047
ADSM Installing the Clients	<i>ADSTAR Distributed Storage Manager: Installing the Clients</i>	SH26-4049
ADSM Client Reference Cards	<i>ADSTAR Distributed Storage Manager: Client Reference Cards</i>	SX26-6013
ADSM Using the Application Programming Interface	<i>ADSTAR Distributed Storage Manager: Using the Application Programming Interface</i>	SH26-4002
ADSM AFS/DFS Backup Clients	<i>ADSTAR Distributed Storage Manager AFS/DFS Backup Clients</i>	SH26-4048
CD-ROM		
ADSM Online Library		SK2T-1893

IBM Redbooks

The International Technical Support Center (ITSC) publishes redbooks, which are 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 E. 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
CD Showcase	Operating System/2	VM/ESA
CICS	Operating System/400	VM/XA
Common User Access	OS/2	VSE/ESA
CUA	OS/400	VTAM
DATABASE 2	POWERparallel	WIN-OS/2
DB2/6000		

The following terms are trademarks of other companies:

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

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.

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*.

availability management. Managing recovery from relatively common computer system outages such as a disk drive head crash. Recovery is often accomplished by using disk mirroring and other forms of RAID technology, or by maintaining onsite backup copies of data.

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.

awk. In AIX, a pattern-matching program for processing text files. With the DRM feature, you can use an awk script to break up the disaster recovery plan file into usable parts.

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.

BACKUPOOL. 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*.

boot media. Media that contains operating system and other files essential to running a workstation or server.

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.

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. 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 or a single client file space 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.

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 Systems Application Architecture 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.

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. The action performed by the DSMSERV DUMPDB utility (DMPADSM command on AS/400), which copies ADSM database entries to media for later reload in case a catastrophic error occurs.

database load. The action performed by the DSMSERV LOADDB utility (LODADSM command on AS/400), which copies ADSM 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*.

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.

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 client. The group of clients supported by ADSM that include clients on OS/2, DOS, Windows, Apple, and Novell NetWare operating systems.

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 some ADSM 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.

disaster recovery. Recovery from catastrophic interruptions of computer systems, such as loss of the system location because of natural events. Backup data is kept offsite to protect against such catastrophes.

Disaster Recovery Manager (DRM). An ADSM feature that assists in preparing and later using a disaster recovery plan for the ADSM server.

disaster recovery plan. A document that contains information about how to recover computer systems if a disaster occurs. With DRM, the plan is a file that contains information about the software and hardware used by the ADSM server, and the location of recovery media.

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.

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.

DRM. Disaster Recovery Manager.

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

dsm.serv.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

GUI. Graphical user interface.

H

HDA. Head-disk assembly.

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

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).

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

machine information. Details about the machine on which a client node resides.

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.

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*.

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 recovery media. Media that is kept at a different location to ensure its safety if a disaster occurs at the primary location of the computer system. The media contains data necessary to recover the ADSM server and clients. The offsite recovery media manager, which is part of DRM, identifies recovery media to be moved offsite and back onsite, and tracks media status.

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 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.

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.

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.

recovery media. Media that contains data necessary to recover the ADSM server and clients.

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. (1) Define a client node or administrator who can access the server. See *registration*. (2) Specify licenses that have been purchased for the server.

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*.

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*.

S

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. 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 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-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). 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.

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 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 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.

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*.

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 45
- activating
 - policy set 45
 - server 17
- add a library to the system 35, 36
- administrative client
 - description xv
 - ending 19
 - name 13
 - name and password 17
 - password 13
 - setting options for 59
 - starting 17
- administrative client icon 7
- administrative command-line interface 53
- administrative graphical user interface (GUI) 53
- administrator
 - granting authority to 23
 - registering 23
- Adobe Acrobat 8
- ADSM books 9
- ADSM server options 46
- ADSM utilities window 7
- ADSM, setting up 17
- advanced tasks xii
- application program interface 70
- applying PTFs 63
- archive
 - description ix, xv, 17
 - to tape 26, 40, 42
- archive copy group ix, 43
- archive pool size 22
- associating
 - client with a policy domain 46
 - schedule with clients 47
- authority level
 - granting 23
- automated library, description 25
- automating
 - administrative commands 50
 - client operations xxi, 47

B

- backup
 - description ix, xv, xvii, 17
 - to tape 26, 40, 42
- backup copy group ix, 43
- backup pool size 22
- backup-archive client
 - description xv
 - register 24
- backup-archive client icon 7
- books, ADSM 9

C

- client communications 59
- client components 1
- client directory 1
- client installation 10
- client node
 - description 17
 - registration 24
- client node password 13
- client options 59
 - setting options for TCP/IP communication 59
- client scheduler icon 7
- client scheduler, starting 49
- collocation 25
- command DSMSERV INSTALL 6
- command-line interface, administrative
 - password 17
 - quitting 54
 - starting 54
 - using 54
- commands, administrative
 - ACTIVATE POLICYSET 45
 - COPY DOMAIN 43
 - DEFINE ASSOCIATION 47
 - DEFINE DEVCLASS 36, 37
 - DEFINE DRIVE 35, 36
 - DEFINE LIBRARY 35, 36
 - DEFINE SCHEDULE 47, 50
 - DEFINE STGPOOL 36, 37
 - DEFINE VOLUME 22
 - DSMSERV DUMPDB 2
 - EXPIRE INVENTORY 46, 50

- commands, administrative (*continued*)
 - HALT 19
 - QUERY EVENT 50, 51
 - QUERY STGPOOL 38
 - REGISTER LICENSE 22
 - REGISTER NODE 24
 - UPDATE COPYGROUP 43
 - UPDATE NODE 46
 - UPDATE STGPOOL 40, 41
 - UPGRADEDDB 3
 - VALIDATE POLICYSET 44
- commands, client
 - QUIT 19, 54
 - SCHEDULE 49
- communication methods
 - NONE 58
 - setting administrative client options for TCP/IP 59
 - setting options for the client 59
 - Shared Memory 59
 - TCP/IP 58
- Compuserve 63
- configuring
 - devices, automated library example 25, 28
 - devices, manual library example 28
- conventions used in this book vii
- COPY DOMAIN command 43
- copy group description xvi
- copying
 - policy domain 43
 - policy set 43
- create additional servers 6
- creating
 - defining 20
 - description xvi
 - disk storage pool volumes 21
 - volumes 20
- customizing
 - adding tape drives 25, 35, 36
 - creating disk storage pool volumes 21, 22
 - initializing a tape volume 38

D

- database volume 6
- database volume, creating 20
- default
 - disk storage pools 21
 - management class 43
 - policy (standard) ix

- DEFINE ASSOCIATION command 47
- DEFINE DEVCLASS command 36, 37
- DEFINE DRIVE command 35, 36
- DEFINE LIBRARY command 35, 36
- DEFINE SCHEDULE command 47, 50
- DEFINE STGPOOL command 36, 37
- define the database 20
- define the recovery log 20
- DEFINE VOLUME command 22, 38
- defining
 - administrator 23
 - archive copy group 43
 - backup copy group 43
 - client node 24
 - device class 25, 35, 36
 - disk storage pool 21
 - drive 35, 36
 - library 25, 35, 36
 - management class 43
 - policy 42, 43
 - storage pool 36, 37
 - storage pool volume 22
- device 25
- device class description xviii, 25
- device driver
 - ADSMSCSI 27, 28, 29
 - for automated tape devices 27
 - for manual tape devices 27, 28, 29
 - for optical devices 28
- device, configuring for ADSM
 - automated tape library 28
 - introduction 25
 - manual tape drive 28
 - optical drive 33
 - overview for removable media 27
- directories, client and server 1
- disaster, protecting against 52
- disk space
 - required for installation 1
- disk storage pools
 - creating 21
 - default 21
 - defining 21
- documentation 8, 9
- drive xviii, 35, 36
- driver, device
 - ADSMSCSI 27, 28, 29
 - for automated tape devices 27
 - for manual tape devices 27, 28, 29
 - for optical devices 28

- dsm.opt file 12
- dsm.sys file 12
- dsmsched.log 50
- DSMSERV DUMPDB command 2
- DSMSERV INSTALL command 6
- dsm serv.dsk file 3, 6, 13
- dump database 2

E

- enabling communications 57
- ending the administrative client 19
- expiration
 - description xviii, 25
 - interval (server option) 46
 - processing xviii, 25, 50
 - server option 17, 46
- EXPIRE INVENTORY command 46, 50

F

- files
 - dsm.opt 12
 - dsm.sys 11
 - dsm serv.dsk 3
 - dsm serv.dsk file 13
 - dsm serv.opt 6
 - dsm serv.opt.smp 57
 - pdf 9
 - server directory 6
 - setgid 10
 - setuid 10

G

- getting started icon 7
- GRANT AUTHORITY command 23
- granting authority to administrators 23
- graphical user interface 17
 - changing an object 53
 - defining an object 53
 - deleting an object 54
 - using 53

H

- HALT command 19
- halting the server 19
- HSM client icon 7

I

- IBM Redbooks 70
- IBMLink 63
- icon
 - administrative client 7
 - backup-archive client 7
 - client scheduler 7
 - getting started 7
 - HSM client 7
 - server 7
- include-exclude list 46
- initializing
 - server 20
 - tape volumes 38
- install ADSM books 9
- install ADSM utilities window 7
- installation results 13
- installation verification 13
- installing
 - ADSM 1
- installing Adobe Acrobat 8
- installing PTFs 63
- installing the administrative client 10
- installing the client
 - directory for client files 11
 - procedure 10
 - update dsm.sys file 11
 - updating dsm.opt file 12
- installing the server
 - minimum requirements 1
 - over version 1 server 2
 - procedure 5
- interactive mode command-line interface 54
- interfaces, administrative
 - command line 53
 - graphical user interface (GUI) 53
- Internet 63

L

- label, tape 38
- library
 - configuration example 28
 - description xviii, 25
 - manual 25, 28
- license, ADSM 22

M

- maintenance address on the Web 63
- maintenance upgrade 63
- management class
 - default 25, 45, 47
 - description xvi, xvii
- MANUAL libraries 28
- manual library, description 25
- migrating to tape 26, 40, 42
- migrating to version 2 2
- migration
 - hierarchy 26, 40
 - server storage xv
 - thresholds 40, 41
 - triggering 41
- minimum requirements for installation 1
- minimum requirements to install the server 1
- mount limit, description 26
- multiple server, create 6

N

- name, administrative client 1, 13
- name, backup-archive client 1
- node
 - See client node, registration

O

- online assistance 63
- online books, viewing 8
- online documentation 9
- optical device support
 - device driver 28
- options file
 - editing 57
- options, client
 - SHMPort 61
 - TCPBuffsize 59, 60
 - TCPPort 58
 - TCPServeraddress 59
 - TCPWindowSize 58
- options, server
 - administrative client, starting 17
 - communications 57
 - initializing ADSM server 20
 - starting the server 17

P

- password, administrative client 13, 17
- password, client node 13
- PDF file names 9
- PDF viewer 8
- planning for ADSM xiii
- policy
 - description xv
 - standard (default) ix
 - using 42, 45
- policy domain
 - associating a client with 43, 46
 - defining 43
- policy set
 - activating 35, 36, 45
 - validating 25, 44
- program temporary fix 63
- protecting data 52
- PTF 63
- publications 9
- publications, viewing 8

Q

- QUERY EVENT command 50, 51
- QUERY STGPOOL command 38
- querying
 - events 50, 51
 - storage pool volumes 38
- QUIT command 54
 - creating 20
 - size 20

R

- recovery log volume 6, 20
- redbooks 70
- REGISTER LICENSE command 22
- REGISTER NODE command 24
- registering
 - administrator 23
 - client node 24
- removing PTFs 64
- requirements for installation 1
- requirements to install the server 1
- results of installation 13

S

- SCHEDULE command 49
- scheduling
 - administrative commands 50
 - client operations 47
- scratch volumes 27
- secondary server attachment license 23
- selective backup 17
- server
 - activating 17
 - authority 23
 - changing options 46
 - creating storage pool volumes 21
 - description xiv, 27
 - halting 19
 - initializing 20
 - installation minimum requirements 1
 - options 46
 - removing from Sun Solaris 15
 - setting up 17
 - starting 17
 - stopping 19
 - storage xiv
 - storage migration xvii, 26, 40
- server communications
 - NONE 58
- server console icon 7
- server directory 1
- server license 22
- server options
 - dsmserv.opt.smp 57
 - tailoring 57
- service 63
- shared memory client options 61
- SHMPort 61
- size
 - archive pool 22
 - backup pool 22
 - space management pool 22
- space management (HSM) xv
- space management pool size xv, 22
- space requirements to install 1
- special file names 34
- standard label 38
- standard policy ix
- start utilities 7
- starting
 - client scheduler 49
 - server 17

- stopping
 - client 19
 - server 19
- storage
 - archive pool size 22
 - ARCHIVEPOOL 21, 26, 40, 42
 - backup pool size 22
 - BACKUPPOOL 21, 26, 40, 42
 - default 42
 - description xiv, xvii, xviii, 17
 - disk 21
 - space management pool size 22
 - SPACEMGPOOL 21, 26, 40, 42
 - volumes xviii
- storage pool 36, 37
- storage pool volumes 6

T

- tape
 - backup to 25
 - drives 25, 35, 36
 - label 38
 - volumes, initializing 38
- tape management system 25
- TCP/IP 59
- TCPBUFFSIZE option 61
- TCPPORT option 58, 61
- TCPSERVERADDRESS option 61
- TCPWINDOWSIZE option 58, 61
- threshold, migration 40
- triggering migration 41

U

- UPDATE COPYGROUP command 43
- UPDATE NODE command 46
- UPDATE STGPOOL command 21, 40
- updating
 - client node 46
 - copy group 43
- upgrade to version 2 2
- UPGRADEDDB command 3
- utilities, ADSM 7
- utilities, start 7

V

- VALIDATE POLICYSET command 44

validating a policy set 25, 44
verify installation 13
version 2, upgrading to 2
view ADSM books 8

W

web information 63

Communicating Your Comments to IBM

ADSTAR Distributed Storage Manager
for Sun** Solaris**
Quick Start
Version 2

Publication No. GC35-0262-00

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 Sun** Solaris**
Quick Start
Version 2**

Publication No. GC35-0262-00

Overall, how satisfied are you with the information in this book?

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

How satisfied are you that the information in this book is:

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Complete	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Easy to find	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Easy to understand	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Well organized	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Applicable to your tasks	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

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
GC35-0262-00



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

GC35-0262-00

Cut or Fold
Along Line



Program Number: 5639-B22



Printed in the United States of America
on recycled paper containing 10%
recovered post-consumer fiber.

GC35-0262-00



Spine information:



ADSTAR Distributed Storage Manager
for Sun** Solaris**

Quick Start

Version 2