

ADSTAR Distributed Storage Manager  
for MVS



# Administrator's Guide

*Version 2*

**Note!**

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

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

### **First Edition (July 1995)**

This edition applies to Version 2 of the ADSTAR Distributed Storage Manager for MVS (5655-119) and to any subsequent releases 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 1993, 1995. 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

<b>Notices</b> . . . . .	xi
Programming Interface . . . . .	xi
Trademarks . . . . .	xiii
<b>Preface</b> . . . . .	xv
Who Should Read This Publication . . . . .	xv
What You Should Know Before Reading This Publication . . . . .	xv
ADSTAR Distributed Storage Manager Publications . . . . .	xv
Related MVS System Publications . . . . .	xvi
Software Developer's Program . . . . .	xvii
Do You Have Comments or Suggestions? . . . . .	xvii
Translations . . . . .	xvii
<b>Changes for ADSTAR Distributed Storage Manager Version 2</b> . . . . .	xix
<b>Chapter 1. Introducing ADSTAR Distributed Storage Manager</b> . . . . .	1
Administrator Tasks . . . . .	2
Using the ADSM Interfaces . . . . .	3
Managing Server Operations . . . . .	3
Managing the Database and Recovery Log . . . . .	3
Managing Licensing, Privilege Classes, and Registration . . . . .	3
Managing Policies . . . . .	4
Scheduling Operations . . . . .	4
Managing Storage Devices . . . . .	4
Managing Storage Pools . . . . .	4
Managing Storage Pool Volumes . . . . .	5
Managing Tape Operations . . . . .	5
Exporting and Importing Data . . . . .	5
Recovering Data . . . . .	5
<b>Chapter 2. Using ADSTAR Distributed Storage Manager Interfaces</b> . . . . .	7
Using the Graphical User Interface . . . . .	8
Getting Started . . . . .	8
Using Online Help . . . . .	15
Using the Administrative Command Line-Interface . . . . .	22
Using Console Mode . . . . .	22
Using Mount Mode . . . . .	23
Using Batch Mode . . . . .	23
Using Interactive Mode . . . . .	23
Using Administrative Client Options . . . . .	24
Using Macros to Issue Commands . . . . .	24
Using the Server Console Session . . . . .	25
Application Programming Interface . . . . .	26
Task, GUI, and Command Cross-Reference . . . . .	27
Tasks Found in "Using ADSTAR Distributed Storage Manager Interfaces" . . . . .	28
Tasks Found in "Managing Server Operations" . . . . .	29
Tasks Found in "Managing the Database and Recovery Log" . . . . .	31
Tasks Found in "Managing Licensing, Privilege Classes, and Registration" . . . . .	33
Tasks Found in "Managing Policies" . . . . .	36
Tasks Found in "Scheduling Operations" . . . . .	39

Tasks Found in “Managing Storage Devices”	41
Tasks Found in “Managing Storage Pools”	42
Tasks Found in “Managing Storage Pool Volumes”	43
Tasks Found in “Managing Tape Operations”	44
Tasks Found in “Exporting and Importing Data”	45
Tasks Found in “Recovering Data”	46
<b>Chapter 3. Managing Server Operations</b>	<b>49</b>
Starting, Halting, and Restarting the Server	49
Starting the Server	49
Halting the Server	50
Restarting the Server	51
Managing Client Sessions	51
Requesting Information about Client Sessions	51
Canceling a Client Session	52
Disabling or Enabling Server Access	53
Managing Server Processes	53
Requesting Information about Server Processes	54
Canceling Server Processes	54
Varying Disk Volumes Online or Offline	55
Requesting Information about Server Status	55
Setting the Server Name	56
Querying Server Options	56
Managing the Activity Log	56
Changing the Size of the Activity Log	57
Setting the Activity Log Retention Period	57
Requesting Information from the Activity Log	58
Diagnosing Problems and Monitoring Errors	58
Monitoring Accounting Records	59
Getting Help on Commands and Error Messages	61
<b>Chapter 4. Managing the Database and Recovery Log</b>	<b>63</b>
Database and Recovery Log	63
Database Information	64
Recovery Log Information	64
Relationship between the Database and the Recovery Log	64
Storage Pool Volume Locations	65
Transaction Processing	65
Allocating Space for the Database and Recovery Log	65
Step 1: Estimating the Amount of the Database Space	66
Step 2: Estimating the Amount of Recovery Log Space	66
Step 3: Understanding How Space is Managed by the Server	67
Step 4: Allocating Space for the Database and Recovery Log	71
Adding Space to the Database or Recovery Log	75
Step 1: Defining Disk Volumes	75
Step 2: Determining the Maximum Extension	76
Step 3: Extending the Capacity of the Database or Recovery Log	77
Deleting Space from the Database or Recovery Log	78
Step 1: Determining the Size of Volumes	79
Step 2: Determining If There Is Sufficient Free Space	81
Step 3: Reducing the Capacity of the Database or Recovery Log	82
Step 4: Deleting a Volume from the Database or Recovery Log	82
Optimizing the Performance of the Database or Recovery Log	83
Database Buffer Pool	84

Recovery Log Buffer Pool . . . . .	84
Monitoring the Database Buffer Pool . . . . .	84
Monitoring the Recovery Log Buffer Pool . . . . .	86
<b>Chapter 5. Managing Licensing, Privilege Classes, and Registration . . .</b>	<b>89</b>
Managing ADSM Licenses . . . . .	89
Licensed Features . . . . .	90
License Compliance . . . . .	92
Licensing Example . . . . .	92
Monitoring Licenses . . . . .	92
Ensuring Client/Server Authentication . . . . .	92
Setting Client Password Authentication . . . . .	93
Setting User Password Expiration . . . . .	93
Registering Administrators or Updating Information . . . . .	93
Granting Administrative Authority . . . . .	94
System Privilege . . . . .	95
Unrestricted Policy Privilege . . . . .	95
Restricted Policy Privilege . . . . .	96
Unrestricted Storage Privilege . . . . .	96
Restricted Storage Privilege . . . . .	97
Operator Privilege . . . . .	97
Analyst Privilege . . . . .	98
Revoking or Reducing Administrative Authority . . . . .	98
Revoking All Administrative Privilege Classes . . . . .	98
Revoking One or More Administrative Privilege Classes . . . . .	99
Reducing System Authority . . . . .	99
Reducing Unrestricted Policy to Restricted Policy Privilege . . . . .	99
Reducing Unrestricted Storage to Restricted Storage Privilege . . . . .	100
Reducing Restricted Policy or Storage Privilege . . . . .	100
Managing Administrator Access . . . . .	100
Renaming an Administrator . . . . .	101
Removing Administrators . . . . .	101
Locking and Unlocking Administrators from the Server . . . . .	101
Managing Client Nodes . . . . .	102
Setting Client Node Registration . . . . .	103
Registering Client Nodes . . . . .	104
Managing Client Node Access . . . . .	105
Requesting Information about Client Nodes or File Spaces . . . . .	107
Deleting File Spaces and Client Nodes . . . . .	109
Registering an Application Programming Interface to the Server . . . . .	110
Understanding How the Compression Option is Set . . . . .	110
Understanding How the File Deletion Option is Set . . . . .	111
<b>Chapter 6. Managing Policies . . . . .</b>	<b>113</b>
Policy Operations . . . . .	113
Backup and Restore . . . . .	114
Archive and Retrieve . . . . .	114
Migration and Recall . . . . .	114
Policy Objects . . . . .	115
Management Classes . . . . .	116
Management Class Configuration . . . . .	117
Default Management Classes . . . . .	117
The Include-Exclude List . . . . .	118
How Files Are Associated with a Management Class . . . . .	118

File Eligibility for Policy Operations	120
Incremental Backup	120
Selective Backup	121
Archive	121
Automatic Migration from a Client Node	121
How Client Migration Works with Backup and Archive	122
Using the Standard Storage Management Policies	122
Creating Your Own Storage Management Policies	123
Defining and Updating a Policy Domain	126
Defining and Updating a Policy Set	127
Defining and Updating a Management Class	127
Defining and Updating a Backup Copy Group	128
Defining and Updating an Archive Copy Group	132
Assigning a Default Management Class	134
Validating and Activating Policy Sets	134
Activating Policy Sets	135
Starting Expiration Processing	135
Querying Policy Objects	136
Querying Copy Groups	136
Querying Management Classes	137
Querying Policy Sets	137
Querying Policy Domains	138
Deleting Policy Objects	138
Deleting Copy Groups	138
Deleting Management Classes	139
Deleting Policy Sets	139
Deleting Policy Domains	139
<b>Chapter 7. Scheduling Operations</b>	<b>141</b>
Central Scheduling	141
Scheduling Modes	142
Client-Polling Scheduling Mode	142
Server-Prompted Scheduling Mode	142
Setting Scheduling Modes	143
Setting the Scheduling Mode on the Server	144
Setting the Scheduling Mode on Client Nodes	144
Coordinating Server and Client Node Schedules	144
Specifying the Schedule Period for Incremental Backup Operations	145
Controlling the Scheduler Workload	145
Controlling Contact with the Server	147
Managing Schedules and Associations	148
Defining or Updating Schedules	149
Copying Schedules	152
Querying Schedules	153
Deleting Schedules	153
Associating Client Nodes with Schedules	153
Querying Associations	154
Deleting Associations	154
Managing Scheduled Event Records	155
Querying Event Records	155
Removing Event Records from the Database	156
<b>Chapter 8. Planning Your Storage Environment</b>	<b>157</b>
How Workstation Files are Stored	157

How the Server Manages Storage	157
Introduction to Devices	158
Introduction to Storage Pools	158
Introduction to Storage Volumes	158
Expanding Your Data Storage Environment	158
Evaluating Your Storage Environment	159
Considering User Needs for Recovering Data	160
<b>Chapter 9. Managing Storage Devices</b>	<b>161</b>
Selecting a Device Type for a Device Class	161
Defining and Updating CARTRIDGE Device Classes	163
Defining FILE Device Classes	168
Requesting Information about a Device Classes	168
Deleting Device Classes	169
<b>Chapter 10. Managing Storage Pools</b>	<b>171</b>
Storage Pools	172
Primary Storage Pools	172
Copy Storage Pools	173
Assigning Volumes in Storage Pools	178
Assigning Random Access Storage Pool Volumes	178
Assigning Sequential Access Storage Pool Volumes	178
Storage Pool Hierarchy	178
Storage Considerations	179
How ADSM Stores Files in a Storage Pool Hierarchy	179
Server Migration of Files	181
Migration Thresholds for Disk Storage Pools	181
Migration Thresholds for Sequential Access Storage Pools	184
Damaged Files	185
The Use of Cache on Disk Storage Pools	185
Considerations for Using Cache	186
Considerations for Not Using Cache	186
Collocation on Sequential Access Storage Pools	186
How the Server Selects Volumes with Collocation Enabled	188
How the Server Selects Volumes with Collocation Disabled	189
Space Reclamation for Sequential Access Storage Pools	190
Reclamation of Offsite Volumes	191
Reclamation Threshold	191
How Collocation Affects Reclamation	191
Reclamation in a Single Drive Library	192
Estimating Space Needs for Storage Pools	192
Estimating Space Needs in Random Access Storage Pools	192
Estimating Space Needs in Sequential Access Storage Pools	194
Defining or Updating Storage Pools	194
Defining a Primary Storage Pool	194
Defining a Copy Storage Pool	197
Using Copy Storage Pools to Improve Data Availability	199
Recreating Damaged Files	200
Backing Up Storage Pools	200
Full Storage Pool Backups	201
Monitoring the Use of Storage Pools	202
Monitoring the Use of Storage Pool Space	202
Monitoring Migration Thresholds	204
Monitoring the Use of Cache Space on Disk Storage	207

Requesting Information on Storage Occupancy . . . . .	208
Deleting a Storage Pool . . . . .	210
Restoring Storage Pools . . . . .	211
<b>Chapter 11. Managing Storage Pool Volumes . . . . .</b>	<b>213</b>
Storage Pool Volumes . . . . .	213
Access Mode for Storage Pool Volumes . . . . .	214
Adding Random Access Storage Pool Volumes . . . . .	215
Defining Storage Pool Volumes . . . . .	216
Updating Storage Pool Volumes . . . . .	216
Monitoring the Use of Storage Pool Volumes . . . . .	217
Requesting General Information about Storage Pool Volumes . . . . .	218
Requesting Detailed Information about Storage Pool Volumes . . . . .	218
Managing Data Storage . . . . .	219
Auditing a Storage Pool Volume . . . . .	220
Requesting Information about Storage Pool Volume Contents . . . . .	223
Moving Files from One Volume to Another Volume . . . . .	225
Moving Data to Other Volumes in the Same Storage Pool . . . . .	225
Moving Data to Another Storage Pool . . . . .	226
Moving Data from an Offsite Volume in a Copy Storage Pool . . . . .	227
Preparing to Move Data . . . . .	227
Requesting Information about the Data Movement Process . . . . .	228
Monitoring the Movement of Data between Volumes . . . . .	228
Deleting Data Storage Pool Volumes . . . . .	228
Deleting a Storage Pool Volume with Data . . . . .	229
Deleting an Empty Storage Pool Volume . . . . .	229
Deleting a Storage Pool Volume That Contains Residual Data . . . . .	230
Restoring Storage Pool Volumes . . . . .	230
<b>Chapter 12. Managing Tape Operations . . . . .</b>	<b>233</b>
Understanding the Server Tape Processing . . . . .	233
Processing Tapes . . . . .	233
Understanding How the Server Selects and Handles Tapes . . . . .	233
Handling Tape Mount Requests . . . . .	235
Querying Mounted Tapes . . . . .	235
Dismounting Tapes . . . . .	235
Setting Up ADSM with Your Tape Management System . . . . .	235
Defining ADSM to Your Tape Management System . . . . .	236
Specifying Permanent Tape Retention . . . . .	236
Setting Up Tape Installation-Wide Exits . . . . .	237
<b>Chapter 13. Exporting and Importing Data . . . . .</b>	<b>241</b>
Data That Can Be Export and Imported . . . . .	241
Preparing to Export or Import Data . . . . .	242
Using Preview before Exporting or Importing Data . . . . .	242
Planning for Sequential Media Used to Export Data . . . . .	243
Monitoring Export and Import Processes . . . . .	244
Requesting Information about an Export or Import Process . . . . .	244
Viewing Information from the Server Console . . . . .	245
Viewing Information from an Administrative Client . . . . .	245
Querying the Activity Log for Export or Import Information . . . . .	246
Exporting Data to Sequential Media Volumes . . . . .	247
Deciding When to Export Data to Tape . . . . .	248
Exporting Server Data . . . . .	249



Exporting Administrator Information . . . . .	249
Exporting Client Node Information . . . . .	250
Exporting Policy Information . . . . .	250
Importing Data from Sequential Media Volumes . . . . .	251
Step 1: Previewing Information before You Import Data . . . . .	252
Step 2: Importing Definitions . . . . .	254
Step 3: Tailoring Data Storage Definitions on the Target Server . . . . .	256
Step 4: Importing File Data Information . . . . .	257
Considerations When Importing Data . . . . .	259
Recovering from Errors during the Import Process . . . . .	259
<b>Chapter 14. Recovering Data . . . . .</b>	<b>261</b>
Database and Recovery Log Mirroring . . . . .	262
Using Mirroring . . . . .	263
Recovering by Using Mirrored Copies of the Database . . . . .	265
Database Backup and Recovery . . . . .	265
Roll-Forward Recovery . . . . .	266
Point-in-Time Recovery . . . . .	266
Determining Which Recovery Mode to Run . . . . .	267
Determining Which Type of Backup to Run . . . . .	267
When to Perform Backups . . . . .	267
Database Backup Trigger . . . . .	268
Volume History Backup Files . . . . .	268
Device Configuration Backup Files . . . . .	269
Using Database Backup Features . . . . .	269
Defining Device Classes for Backup . . . . .	269
Establishing Volume History Backup Files . . . . .	270
Establishing Device Configuration Backup Files . . . . .	271
Performing the Database Backups . . . . .	273
Setting the Database Backup Trigger . . . . .	273
Setting the Recovery Log Mode . . . . .	275
Adjusting the Size of the Recovery Log . . . . .	275
Recovering by Using Backed Up Copies of the Database . . . . .	276
Restoring a Database by Using Point-in-Time Recovery . . . . .	276
Restoring a Database to its Most Current State . . . . .	278
Storage Pool Backup and Recovery . . . . .	279
Using Storage Pool Backup Features . . . . .	280
Recovering by Using Backed Up Copies of Storage Pools . . . . .	280
ADSM Backup and Recovery Scenarios . . . . .	281
Backup . . . . .	281
Recovery from a Disaster . . . . .	282
Recovery from Media Loss . . . . .	284
Database Salvage Utilities . . . . .	284
Examples of Recovering a Corrupted Database Using DSMSEV Salvage Utilities . . . . .	285
<b>Glossary . . . . .</b>	<b>287</b>
<b>Index . . . . .</b>	<b>303</b>



---

## 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. Refer to the HONE SALESMANUAL or product announcement letters for the most current product information.

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.

---

## Programming Interface

This publication is intended to help the customer plan for and manage the ADSM server.

This publication also documents General-use Programming Interface and Associated Guidance Information, Product-sensitive Programming Interface and Associated Guidance Information, and Diagnosis, Modification or Tuning Information provided by ADSM.

General-use programming interfaces allow the customer to write programs that obtain the services of ADSM.

General-use Programming Interface and Associated Guidance Information is identified where it occurs, either by an introductory statement to a chapter or section or by the following marking:

\_\_\_\_\_ General-use programming interface \_\_\_\_\_

General-use Programming Interface and Associated Guidance Information...

\_\_\_\_\_ End of General-use programming interface \_\_\_\_\_

Product-sensitive programming interfaces allow the customer installation to perform such tasks as diagnosing, modifying, monitoring, repairing, tailoring, or tuning of ADSM. Use of such interfaces creates dependencies on the detailed design or implementation of the IBM software product. Product-sensitive programming interfaces should be used only for these specialized purposes. Because of their dependencies on detailed design and implementation, it is to be expected that programs written to such interfaces may need to be changed in order to run with new product releases or versions, or as a result of service.

Product-sensitive Programming Interface and Associated Guidance Information is identified where it occurs, either by an introductory statement to a chapter or section or by the following marking:

Product-sensitive programming interface

Product-sensitive Programming Interface and Associated Guidance Information...

End of Product-sensitive programming interface

Diagnosis, Modification or Tuning Information is provided to help the customer to do diagnosis of ADSM.

**Warning:** Do not use this Diagnosis, Modification or Tuning Information as a programming interface.

Diagnosis, Modification or Tuning Information is identified where it occurs, either by an introductory statement to a chapter or section or by the following marking:

Diagnosis, Modification or Tuning Information

Diagnosis, Modification or Tuning Information...

End of Diagnosis, Modification or Tuning Information

---

## Trademarks

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

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

The following terms are trademarks of other companies:

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

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.

Windows is a trademark of Microsoft Corporation.

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



---

## Preface

ADSTAR Distributed Storage Manager (ADSM) is a client/server program that provides storage management to customers in a multivendor computer environment. ADSM provides automated, centrally-scheduled, policy-managed backup, archive, and space management services for file-servers and workstations. The Hierarchical Storage Management (HSM) client is the portion of ADSM that provides space management services.

---

### Who Should Read This Publication

This reference is intended for anyone who has been assigned an administrative privilege class. While ADSM can be managed by a single administrator, administrative responsibilities can be divided among a number of people as an installation requires.

All of the administrator commands you need to operate and maintain ADSM can be invoked from a workstation connected to the server.

---

### What You Should Know Before Reading This Publication

You should be familiar with the operating system on which the server resides and the communication protocols required for the client/server environment.

For information on product requirements for ADSM, refer to *ADSM Licensed Program Specifications*.

For information on installing ADSM, refer to *ADSM Installing the Server and Administrative Client*.

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 random access media and sequential access media.

---

### ADSTAR Distributed Storage Manager Publications

The following table lists ADSM publications.

The ADSM library is also available in softcopy on the *ADSM Online Product Library* CD-ROM (order number SK2T-8714).

Short Title	Publication Title	Order Number
ADSM General Information	<i>ADSTAR Distributed Storage Manager: General Information</i>	GH35-0131
ADSM Messages	<i>ADSTAR Distributed Storage Manager: Messages</i>	SH35-0133
ADSM Device Configuration	<i>ADSTAR Distributed Storage Manager: Device Configuration</i>	SH35-0137

<b>Short Title</b>	<b>Publication Title</b>	<b>Order Number</b>
ADSM Licensed Program Specifications	<i>ADSTAR Distributed Storage Manager for MVS: Licensed Program Specifications</i>	GH26-4038
ADSM Installing the Server and Administrative Client	<i>ADSTAR Distributed Storage Manager for MVS: Installing the Server and Administrative Client</i>	SH26-4043
ADSM Administrator's Reference	<i>ADSTAR Distributed Storage Manager for MVS: Administrator's Reference</i>	SH26-4040
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

## Related MVS System Publications

The following table lists publication titles and order numbers for related MVS publications.

<b>Short Title</b>	<b>Publication Title</b>	<b>Order Number</b>
General Information for System Product Version 4	<i>MVS/ESA General Information for MVS/ESA System Product Version 4</i>	GC28-1600
DFSMS/MVS 1.1 General Information	<i>DFSMS/MVS Version 1 Release 1 General Information</i>	GC26-4900



---

## 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 programming 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 *ADSM General Information*. Contact your IBM representative for more information about the translated publications and whether these translations are available in your country.



---

# Changes for ADSTAR Distributed Storage Manager Version 2

The new functions for ADSM Version 2 are:

## **Database backup and recovery**

You can perform full and incremental backups of the server database to protect against loss or damage. You can use the backup copies to restore the database to its current state or to a specific point in time. You can back up the database while the server is available to clients.

**Note:** To allow for recovery of the database to its most current state, you may have to extend your recovery log space significantly.

See Chapter 14, "Recovering Data" on page 261 for details.

## **Storage pool backup and recovery**

You can back up client files stored on storage pools to sequential media. These media can be either onsite, to protect against media loss, or offsite, for disaster recovery purposes.

See Chapter 10, "Managing Storage Pools" on page 171 for details.

## **Administrative command scheduling**

You can define schedules for automatically issuing administrative commands once or periodically.

See Chapter 7, "Scheduling Operations" on page 141 for details.

## **Hierarchical storage management**

Hierarchical storage management (HSM) provides space management services to HSM clients. HSM clients can automatically migrate user files to storage pools to free up client storage space. A user can access a migrated file as if it were on local storage.

See Chapter 6, "Managing Policies" on page 113 for details.

## **Device support enhancements**

Device class types now include support for files. The FILE device type is used for special device classes whose storage volumes are not physical units, such as tape cartridges. Instead, data is written sequentially into standard files in the file system of the server machine.

See Chapter 9, "Managing Storage Devices" on page 161 for details.



---

# Chapter 1. Introducing ADSTAR Distributed Storage Manager

ADSTAR Distributed Storage Manager (ADSM) is an enterprise-wide storage management application for the network. It provides automatic storage management services to multivendor workstations, personal computers, and local area network (LAN) file servers. ADSM includes the following components:

## **Server**

Allows a host system to provide backup, archive, and space management services to workstations. The server maintains a database and recovery log for ADSM resources, users, and user data.

The server controls the ADSM data storage, or storage pools. These are groups of random and sequential access media that store files backed up, archived, and migrated from client nodes.

## **Administrative client**

Allows administrators to control and monitor server activities, define management policies for workstation files, and set up schedules to provide services at regular intervals.

## **Backup-archive client**

Allows users to register their workstations with a server as client nodes. ADSM users can maintain backup versions of their files, which they can restore if the original files are lost or damaged. Users can also archive files that they do not currently need on their workstations and retrieve the archived files when necessary.

## **Hierarchical storage management client**

Provides space management services for workstations. ADSM users can free up workstation storage by migrating less frequently used files to server storage pools. These migrated files are also called *space-managed files*. Users can recall space-managed files automatically simply by accessing them as they normally would.

## **Application programming interface (API)**

Allows users to enhance existing applications with back up, archive, restore, and retrieve services. When users install the ADSM application client on their workstations, they can register as client nodes with an ADSM server.

Figure 1 on page 2 shows an example of an ADSM client/server environment. In this example, an administrator monitors the system from a workstation on which the administrative client program has been installed.

The backup-archive client program and hierarchical client program have been installed on workstations connected through a LAN and registered as client nodes. From these client nodes, users can back up, archive, or migrate files to the server.

Based on ADSM policies assigned to files, the server stores workstation files on disk or tape volumes in data storage, which can be grouped into storage pools.

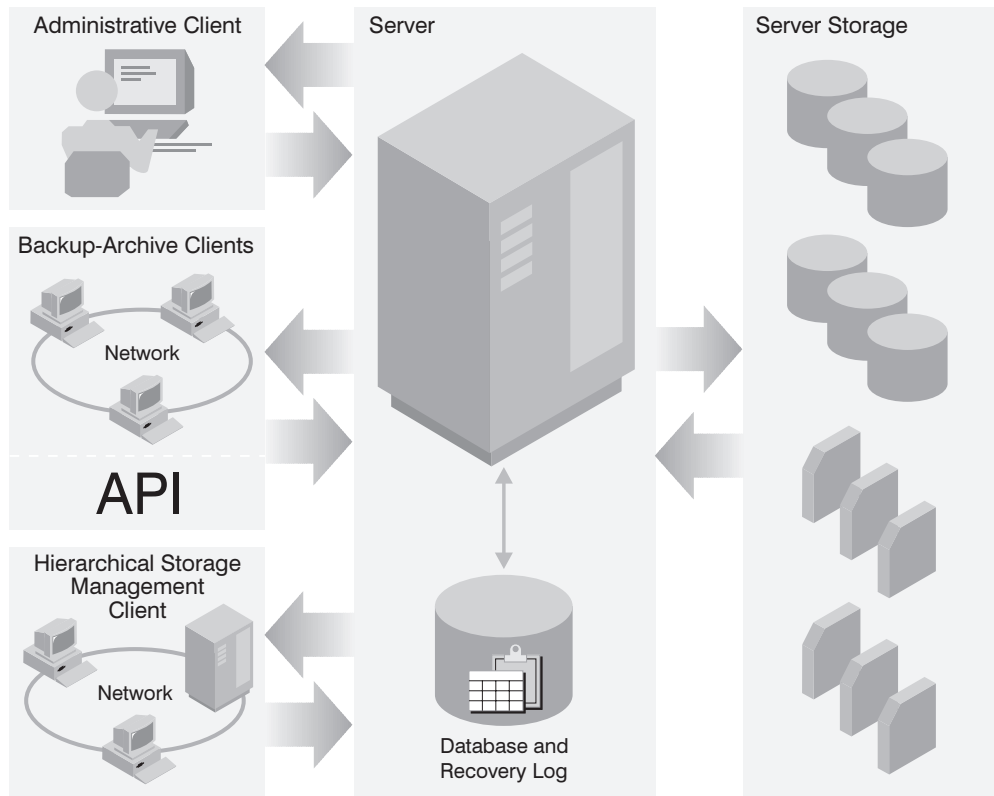


Figure 1. Sample Client/Server Environment

## Administrator Tasks

This section provides a brief overview of the tasks that ADSM administrators can do. It also points to the sections in this publication that present the details of those tasks and the concepts you need to understand to complete them. This section presents the tasks in the order in which they appear in the chapters of this book:

- Using the ADSM interfaces
- Managing server operations
- Managing the database and recovery log
- Managing licensing, privilege classes, and registration
- Managing ADSM policies
- Scheduling ADSM operations
- Managing storage devices
- Managing storage pools
- Managing storage volumes
- Managing tape operations
- Using tape installation-wide exits
- Exporting and importing data
- Providing disaster recovery

## Using the ADSM Interfaces

There are three interfaces to ADSM:

- Graphical user interface (GUI)
- Command-line interface
- Application programming interface

For information about these interfaces and for tables that relate GUI locations and administrative commands with specific tasks, see Chapter 2, “Using ADSTAR Distributed Storage Manager Interfaces” on page 7.

## Managing Server Operations

You can manage server operations such as starting and stopping the server, maintaining and suspending client sessions with the server, and controlling server processes.

ADSM provides you with many sources of information about server and client status and activity, the state of the database, and resource usage. By monitoring this information, you can provide reliable services to users while making the best use of available resources.

For details about the day-to-day tasks involved in administering the server and about reports and information available to you, see Chapter 3, “Managing Server Operations” on page 49.

## Managing the Database and Recovery Log

The ADSM database contains information about the client data in storage pools, registered client nodes, and ADSM policies. The server recovery log, which records changes made to the database, is used to restore the database to a consistent state.

You manage the database and recovery log space and the buffer pool to tune database and recovery log performance.

For more information about the ADSM database and recovery log and about the tasks associated with administering them, see Chapter 4, “Managing the Database and Recovery Log” on page 63.

## Managing Licensing, Privilege Classes, and Registration

You can monitor an installation’s compliance with the terms of its license agreement. ADSM lets you check license compliance and modify the terms.

An organization may name a single administrator or may distribute the workload among a number of administrators and grant them different levels of authority.

You register workstations as client nodes with the server. You can also provide client/server authentication by requiring the use of passwords to ensure that the client and the server are authorized to communicate with each other.

For more information about the preceding concepts and tasks, see Chapter 5, “Managing Licensing, Privilege Classes, and Registration” on page 89.

## Managing Policies

From a client node, users can back up or archive files to the server. This process ensures that current data can be restored or retrieved if it is accidentally deleted or corrupted on their workstations. Users can also migrate files to server storage. This process frees up space on their workstations, but the users can recall the files when they are needed.

You define policies based on user requirements for backing up, archiving, or migrating data. You do this by defining policy objects, which identify backup, archive, and migration criteria, and by scheduling client operations.

For more information about establishing and managing policies for your organization, see Chapter 6, “Managing Policies” on page 113.

## Scheduling Operations

You can define schedules for the automatic processing of most administrative commands and client operations such as backup and restore.

For more information about scheduling ADSM commands and operations, see Chapter 7, “Scheduling Operations” on page 141.

## Managing Storage Devices

A device class represents a set of storage devices with similar availability, performance, and storage characteristics. You must define device classes for the drives available to an ADSM server. You specify a device class each time you define a storage pool, which is a named collection of volumes for storing user data.

For more information about defining device classes, see Chapter 9, “Managing Storage Devices” on page 161.

## Managing Storage Pools

Backed up, archived, or space-managed files are stored in groups of volumes called storage pools. The data on these primary storage pools can be backed to copy storage pools. Because each storage pool is assigned to a device class, you can logically group your storage devices to meet your storage management needs.

You can establish a hierarchy of storage pools. The hierarchy may be based on the speed or the cost of the devices associated with the pools. ADSM migrates user files through this hierarchy to ensure the most efficient use of a server's storage devices.

When defining or modifying a storage pool, you can specify any or all of the following:

- Cache** When files are migrated from disk storage pools, duplicate copies of the files may remain in cache (disk storage) for faster retrieval and are deleted only when space is needed.
- Collocation** ADSM keeps each user's files on a minimal number of volumes within a storage pool. Because user files are consolidated, restoring collocated files requires fewer media mounts.



**Reclamation** Files on sequential access volumes may expire, move, or be deleted. The reclamation process consolidates the active, unexpired data on many volumes onto fewer volumes. The original volumes can then be reused for new data.

For more information about understanding and defining storage pools and taking advantage of storage pool features, see Chapter 10, “Managing Storage Pools” on page 171.

## Managing Storage Pool Volumes

You manage storage volumes not only by defining, updating, and deleting volumes, but also by monitoring the use of server storage. Monitoring volumes can reveal inconsistencies that can be corrected between information in the database and client node files in storage pools. You can also move files within and across storage pools to optimize the use of server storage.

For more information about these tasks, see Chapter 11, “Managing Storage Pool Volumes” on page 213.

## Managing Tape Operations

ADSM provides commands and options to help you manage how tapes are used. You can:

- Control tape mount request messages
- Display information about the status of mounted tapes
- Dismount tapes

You can also set up ADSM to work with your tape management system. ADSM provides a deletion installation-wide exit to notify your tape management system when the server deletes a tape from the database.

For more information about these tasks, see Chapter 12, “Managing Tape Operations” on page 233.

## Exporting and Importing Data

As your storage needs increase, you can move data from one server to another. This process is accomplished by exporting part or all of a server’s data to tape or a flat file so that you can then import the data to another server.

For more information about importing data between servers, see Chapter 13, “Exporting and Importing Data” on page 241.

## Recovering Data

ADSM provides a number of ways to recover from media failure or from the a loss of database or storage pools due to a disaster. These recovery methods are based on the following measures:

- Mirroring, by which the server maintains one or more copies of the database or recovery log, allowing the system to continue when one of the mirrored disks fails
- Periodic backup of the database
- Periodic backup of the storage pools
- Offline dump of the database

For more information about preparing for a disaster and for details about recovering from a disaster, see Chapter 14, “Recovering Data” on page 261.

---

## Chapter 2. Using ADSTAR Distributed Storage Manager Interfaces

Users have three interfaces to ADSM:

- Graphical user interface (GUI)
- Command-line interface
- Application programming interface (API)

This chapter describes those interfaces. It also describes how to use the GUI and the command-line interface.

The sections listed in the following table begin at the indicated pages.

Section	Page
<b>Tasks:</b>	
Getting started with the GUI	8
Using online help	15
Using the command-line interface in console mode	22
Using the command-line interface in mount mode	23
Using the command-line interface in batch mode	23
Using the command-line interface in interactive mode	23
Using administrative client options	24
Using macros to issue commands	24
Using the server console session	25

See “Task, GUI, and Command Cross-Reference” on page 27 for task, GUI, and command cross-referencing.

---

## Using the Graphical User Interface

You can manage the server from a graphical user interface on a workstation. This section contains information about:

- Getting started
- Using online help

## Getting Started

This section contains information about:

- Starting the graphical user interface
- Using menu choices
- Completing a task by using the interface
- Closing an administrative client session

## Making Selections by Using the Mouse or the Keyboard

You can make a selection on the graphical user interface by using your mouse or your keyboard:

<b>Mouse</b>	Hold down the left mouse button and drag the cursor over the area to be selected.
<b>Keyboard</b>	Press the control key and select an object with the space bar.

**Starting the Graphical User Interface:** Start the graphical user interface by entering the product command from a window.

## 1

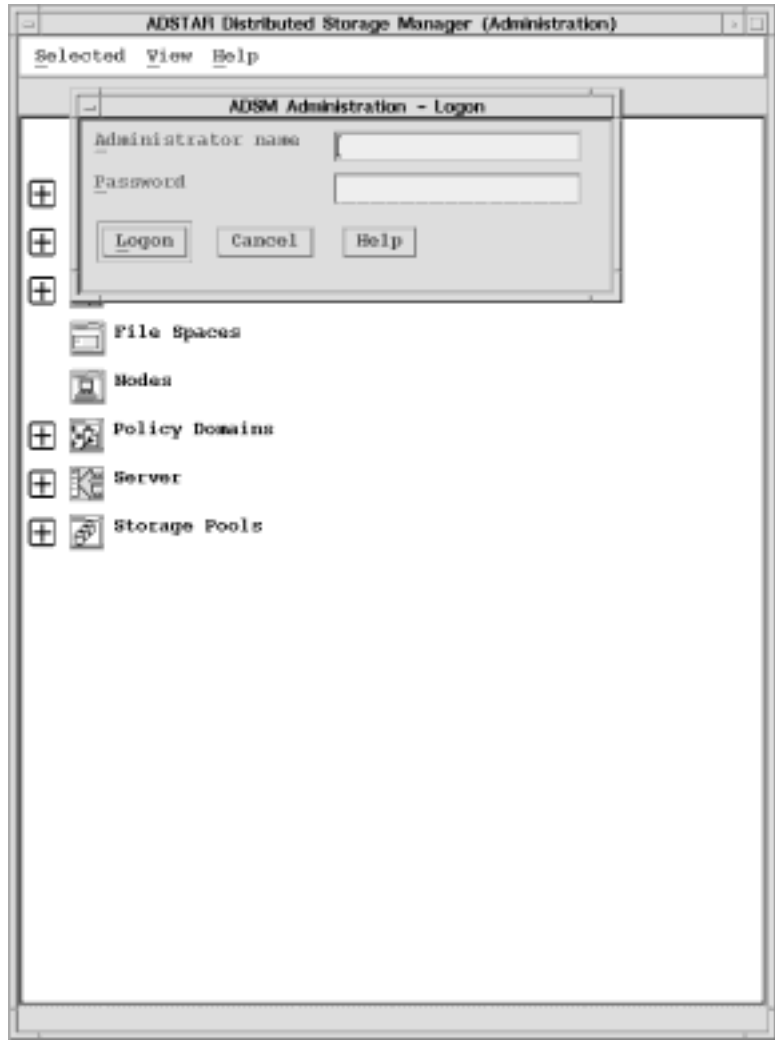
Enter **dsmadm** at the command prompt.

## 2

Type the administrator name and password in the Logon window and then select **Logon** or press the Enter key.

**Note:** If server authentication is off, you do not need to type the password.

The ADSM administrative objects/icons appear.



Your administrative privilege determines which ADSM tasks you can perform and which objects you will work with most frequently.

<b>Administrative Privilege</b>	<b>Objects</b>
<b>Analyst</b>	Server
<b>Operator</b>	Server
<b>Policy</b>	Policy, Nodes, Central Scheduler, and File Spaces
<b>Storage</b>	Data Storage, Database and Recovery Log
<b>System</b>	All

You can run multiple instances of the interface. Each instance of the interface is a new administrative session. Therefore, you can log on to more than one server from the same administrative client by invoking more than one instance of the graphical user interface.

### Using Menu Choices

In the graphical user interface, each window contains menu bars. Each choice in the menu bar leads to an associated pull-down menu. Choices in the pull-down menus are actions or routing choices that relate to the contents of the window. A routing choice with an ellipsis opens another window. A routing choice with an arrow displays a cascaded menu. Objects in the icon pull-down view of the graphical user interface contain some of the following menu choices:

<b>Menu</b>	<b>Contents</b>
<b>Selected</b>	Contains a number of options, for example: Open as: <ul style="list-style-type: none"><li>• Icon</li><li>• Details</li><li>• Properties</li></ul>
<b>Edit</b>	Contains choices that are standard across many different types of objects such as choices that allow you to add new objects or delete existing ones.
<b>View</b>	Contains choices that affect the way an object is displayed.
<b>Help</b>	Contains choices that provide access to help information. You can select a help index, general help, instructions about how to use help, task help, or product information.

**Completing a Task by Using the Interface:** The following example describes how to use the graphical user interface to view and change information about a node. An ADSM administrator with any privilege class can view information about nodes registered with ADSM. In this scenario, you will be displaying information about the **Nodes** object in a domain.

To open the **Icons** view, for example, of the **Nodes** object:

**1**

Select the **Nodes** icon.

**2**

Select **Selected** from the menu bar. Then select **Open as**.

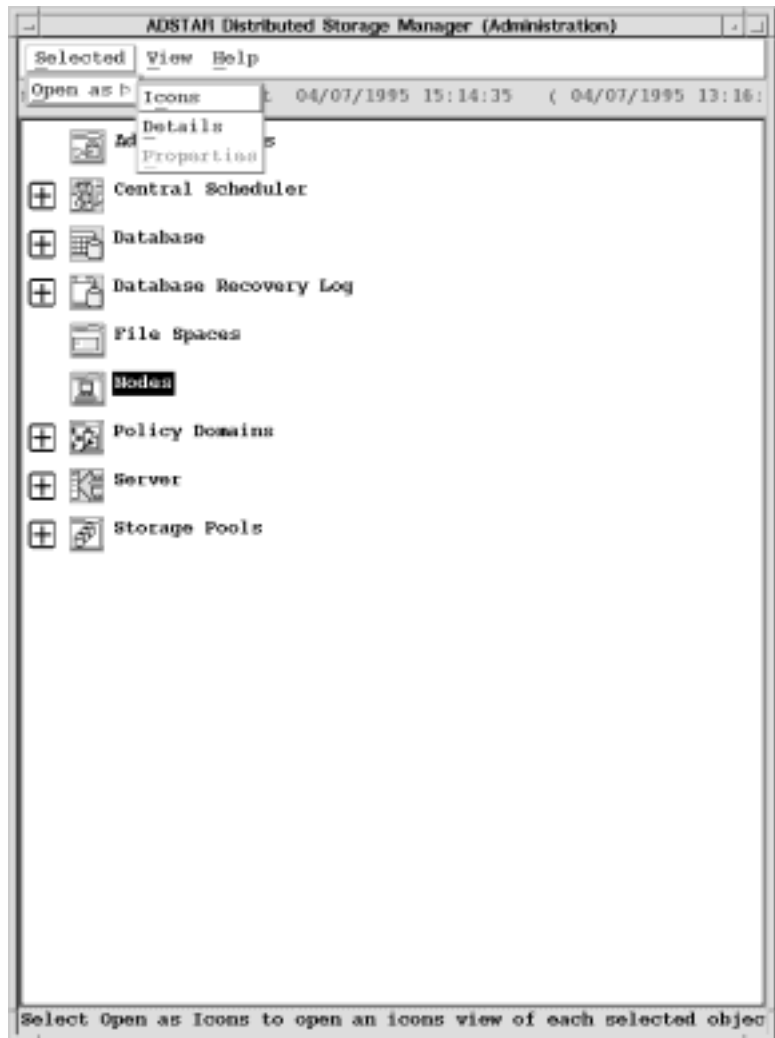
**3**

Select **Icons** from the pull-down menu.

The **Nodes - Icon** screen appears.

**4**

To close the session, see "Closing an Administrative Client Session" on page 14.



The Details view displays information about the nodes. This information includes the name of the node, its operating system (which the node uses to access the server), the policy domain to which the node is assigned, the number of days that have elapsed since the node last accessed the server, and the number of days that have elapsed since a new password was set for the node. In addition, you can see if the node is locked from accessing the server. Move the horizontal and vertical scroll bars to view all of this information.

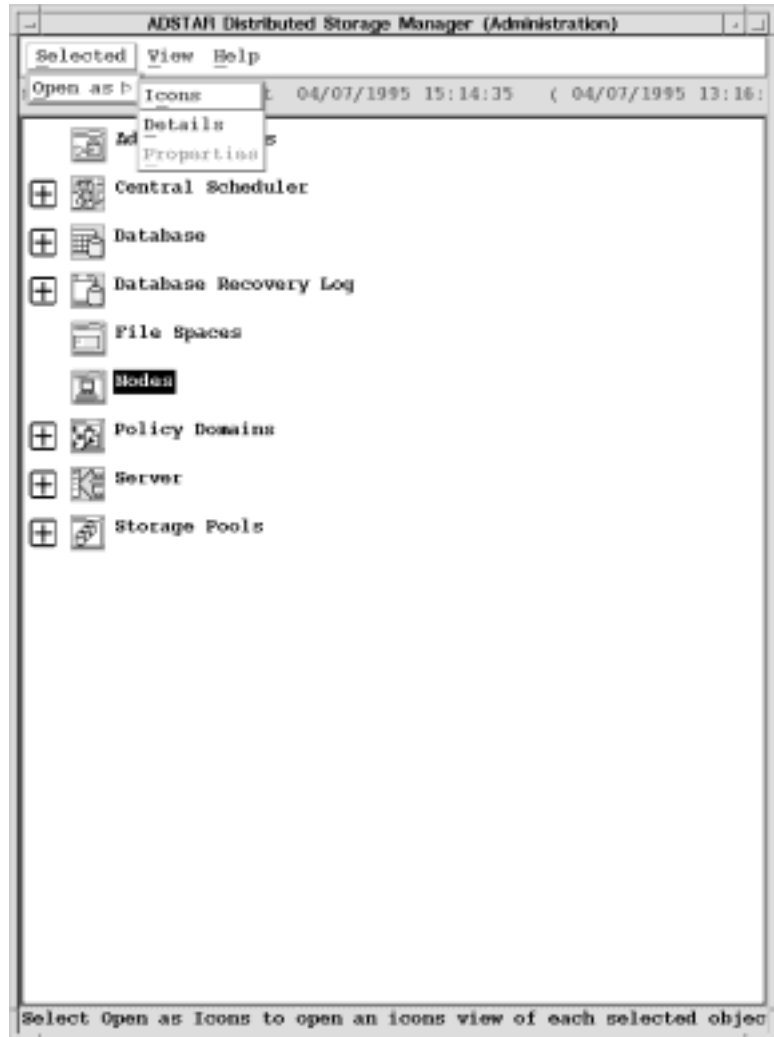
For example, to open the **Details** view window of the **Nodes** object:

### 1

Select the **Nodes** icon.

### 2

Select **Selected** from the menu bar and then select **Open as**.



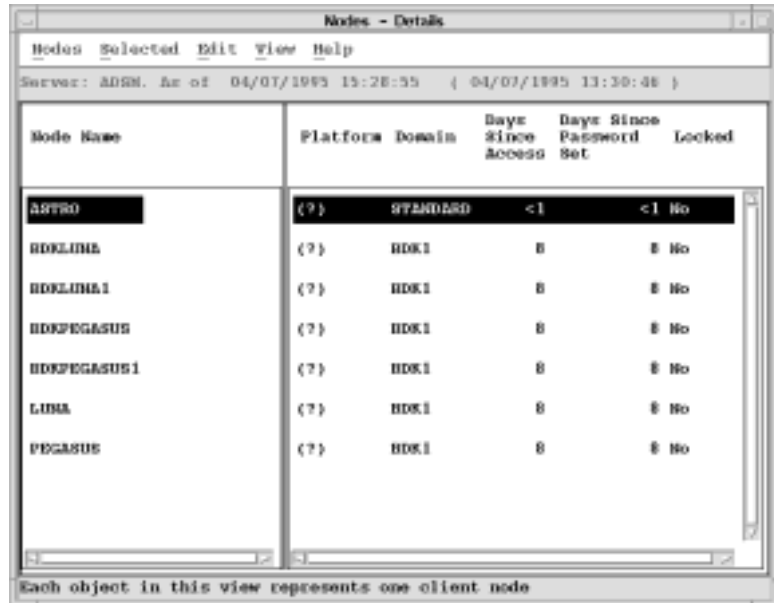
### 3

Select **Details** from the pull-down menu.

The **Nodes - Details** screen appears.

### 4

To close the session, see “Closing an Administrative Client Session” on page 14.





The Properties view window of a Node uses the following notebook tabs to display information about the objects/icon:

- The **General** notebook tab displays the node name, contact information, policy domain, file compression status, and locked status on the first page. The second page of the **General** notebook section shows platform access and registration information.
- The **Password** notebook tab allows you to change the password.
- The **Filespace** notebook tab displays whether the node owner can delete the node's backed up or archived files.
- The **Accounting** notebook tab displays valuable accounting information.

If you have policy or system privilege, you can change information in the Properties view.

For example, to open the **Properties** view window of the **Nodes** object, and to see additional information or change information:

## 1

Select the **Nodes** icon.

## 2

Select **Selected** from the menu bar and then select **Open as**.

## 3

Select **Icons** from the pull-down menu.

## 4

Select a node from the Nodes-Icons screen.

## 5

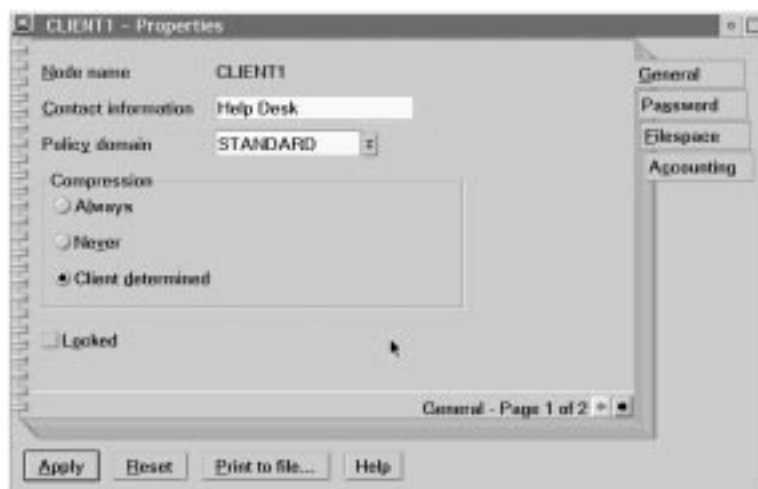
Select **Selected** from the menu.

## 6

Select **Open as properties**.

The **Nodes - Properties** screen appears.

Use the tabs and arrow buttons at the bottom to navigate through the pages of the notebook.



To change information in an entry field of the **Properties** view window:

## 1

Select the field and enter the new information.

## 2

Press the Enter key to accept the information.

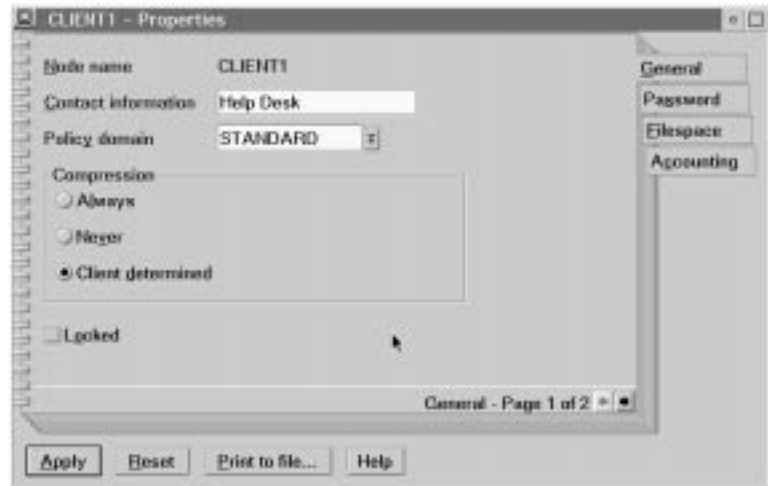
## 3

By using the mouse or cursor keys, change the status of the radio buttons, check button or select from the items in a drop down list.

When you have finished editing the data, select the **Apply** push button.

## 4

To close the session, see “Closing an Administrative Client Session” on page 14.



## Closing an Administrative Client Session

You can close an administrative client session in any one of the following ways:

- Double-click on the system menu symbol.
- Select the system menu. Then select **Close** from the pull-down menu that appears.
- Press Alt+F4 to close the most recent application window.

## Using Online Help

This section describes how to use online help to assist you with the following information:

- Understanding the types of help available
- Accessing help from windows
- Multiple ways to access task help
- Accessing contextual help
- Finding information in the help facility

### Understanding the Types of Help Available

You can request help to receive information about tasks, windows, and the selectable objects, menu choices, notebook tabs, fields, controls or push buttons. All of the following, except contextual help, are available from the Help menu.

The ADSM administrative graphical user interface provides the following types of online help:

<b>Task help</b>	Provides a list of tasks that can be completed with the selected object. When you select a task, the help provides step-by-step information on how to complete the task.
<b>General help</b>	Provides an overview of the function of the window. For example, general help for the Nodes—Icons window explains how to use the window, lists the tasks you can complete with the Nodes object, and displays links (routing choices) where you can obtain more details or related information about the Node object.
<b>Contextual help</b>	Provides specific information for each selectable object, menu choice, notebook tab, field, and control or push button in a window. For example, you can access help for which information to enter in the password field.
<b>Help index</b>	Provides an alphabetic listing of all the help topics. For example, within the help index, you can select <b>Keys Help</b> for information on how to use function keys or short cut keys in the ADSM administrative graphical user interface.
<b>Using help</b>	Provides information on how to access the different types of help.
<b>Product information</b>	Provides ADSM product information, including the copyright notice and release information.

**Accessing Help from Windows:** To access help in an ADSM administrative graphical user interface:

- Select **Help** from the menu bar
- Select an object/icon and press F1
- Select the Help button, if available, to display contextual help on a field or control that has cursor focus
- Press F2 from any contextual or task help window to display **General help**

**Accessing General Help:** You can access general help for a window from the help index, from the help table of contents, from the help menu within the graphical user interface, or from a contextual or task help window of an application window:

**From the help index:**

In the Help **index**, select the name of the general help window that you want to view.

**From the help table of contents:**

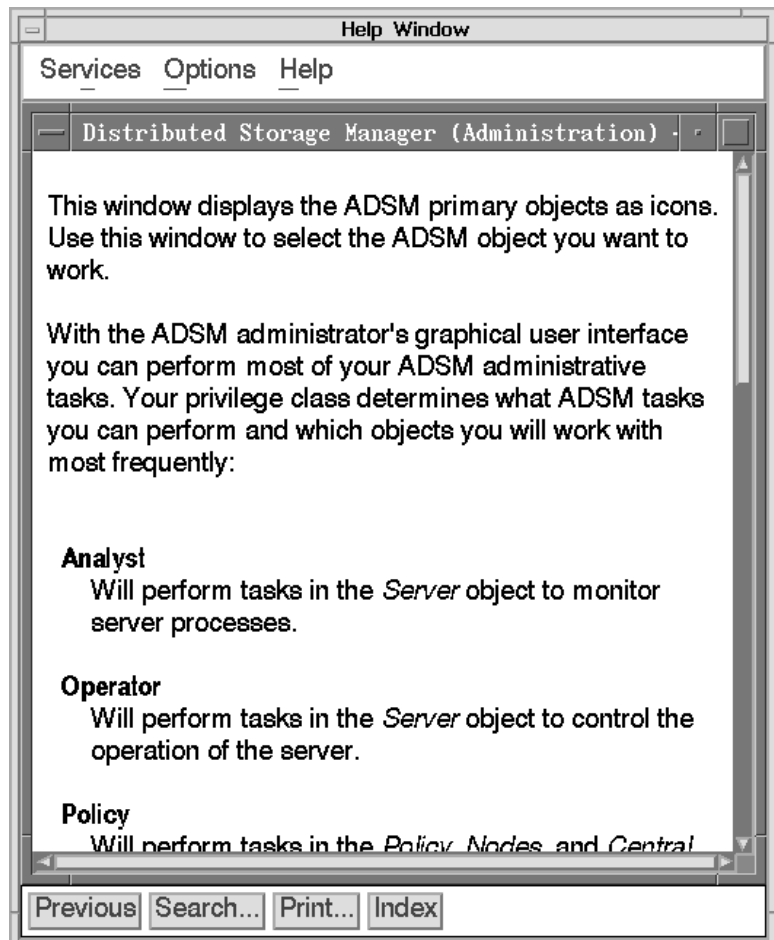
In the Help system, select **Contents** from the **Options** menu and select the help title from the table of contents list.

**From the help menu:**

Select **General help** from the pull-down menu under **Help** on the menu bar.

**From a contextual or task help window:**

Press F2.



**Multiple Ways to Access Task Help:** To view **Task help** about the **Nodes** object/icon, for example, select one of the following:

- **Task help** from the help menu
- The task you need help with from the **Help index**
- Related Information from the **General help** window

To access **Task help**:

**1**

Select **Nodes**, for example.

**2**

Select **Help** from the menu bar on the Nodes Window.

**3**

Select **Task help** from the pull-down menu.

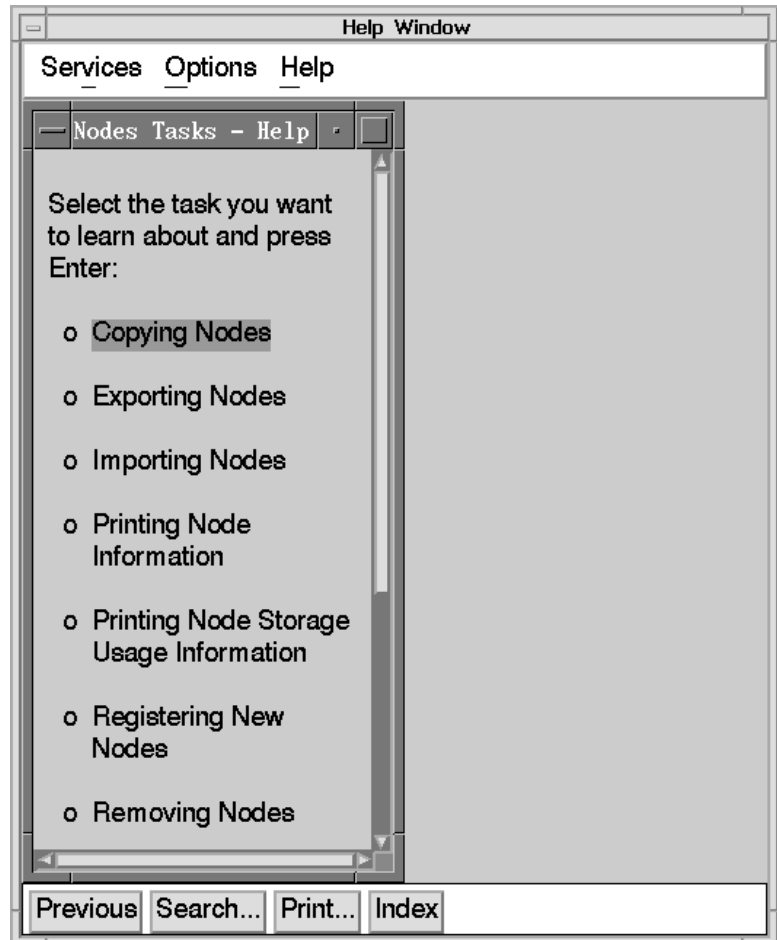
**4**

Follow the instructions in the windows for procedural information for a particular task.

A split window appears. Select a task from the task list on the left. The steps for the task display on the right.

**5**

Close the help window.



**Accessing Contextual Help:** You can access contextual help from a window by moving the cursor to a selectable object, menu choice, notebook tab, field, control, or push button.

To select contextual help:

**1**

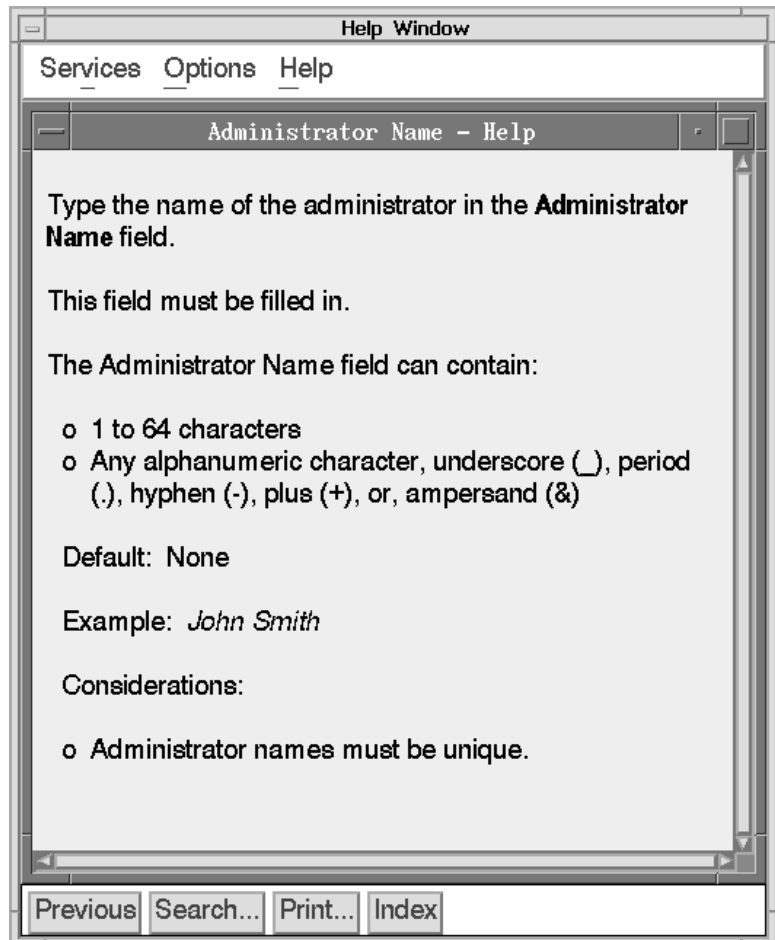
From the Logon window select, for example, the **Password** field.

**2**

Press F1.

**OR**

Select the **Help** push button, if it is displayed on the window.



**Finding Information in the Help Facility:** After you have accessed the help facility, you can use the following help features:

- Hypertext links
- Search facility

**Using Hypertext Links:** Certain words or phrases, known as hypertext links, appear in a different color from the rest of the help text. Select the hypertext links to display related help:

**1**

Select the **Administrators** icon, for example.

**2**

Press F1.

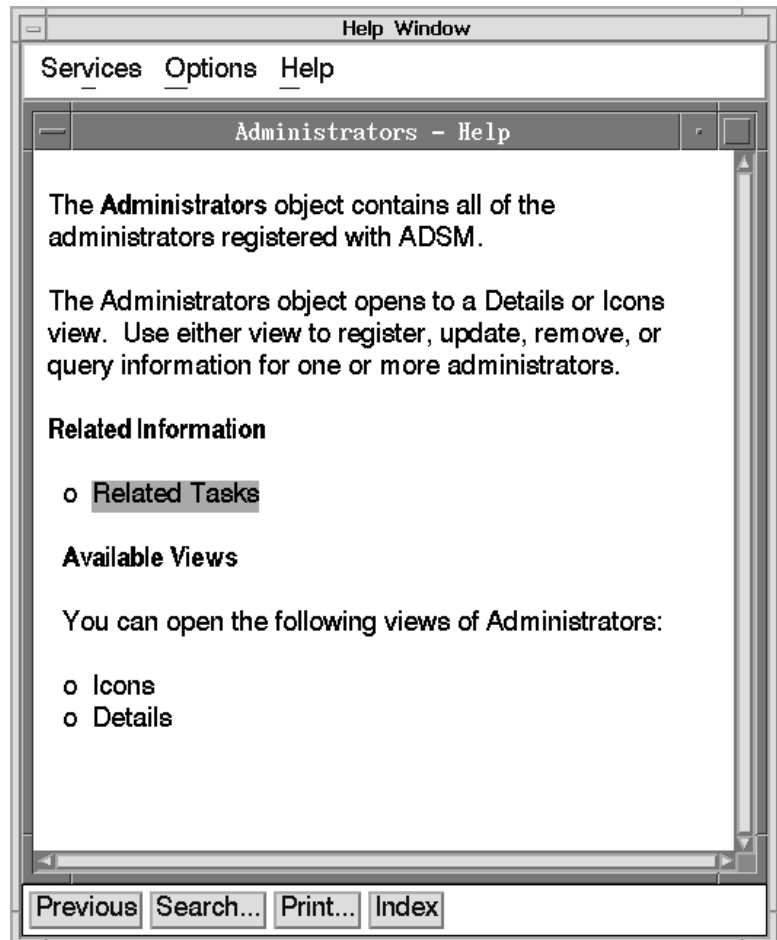
For additional help information:

**1**

Select **Related Tasks**, **Icons**, or **Details**.

**2**

To close the window, use the Escape key or select the **Previous** push button.



**Using the Search Facility:** You can use the search facility in the ADSM administrative graphical user interface help to find occurrences of words or phrases.

To access the search facility on a help panel:

**1**

Select **Administrators** object, for example.

**2**

Select **General help** from the help menu.

**3**

Select **Search** from the Services menu, or select the Search push button.

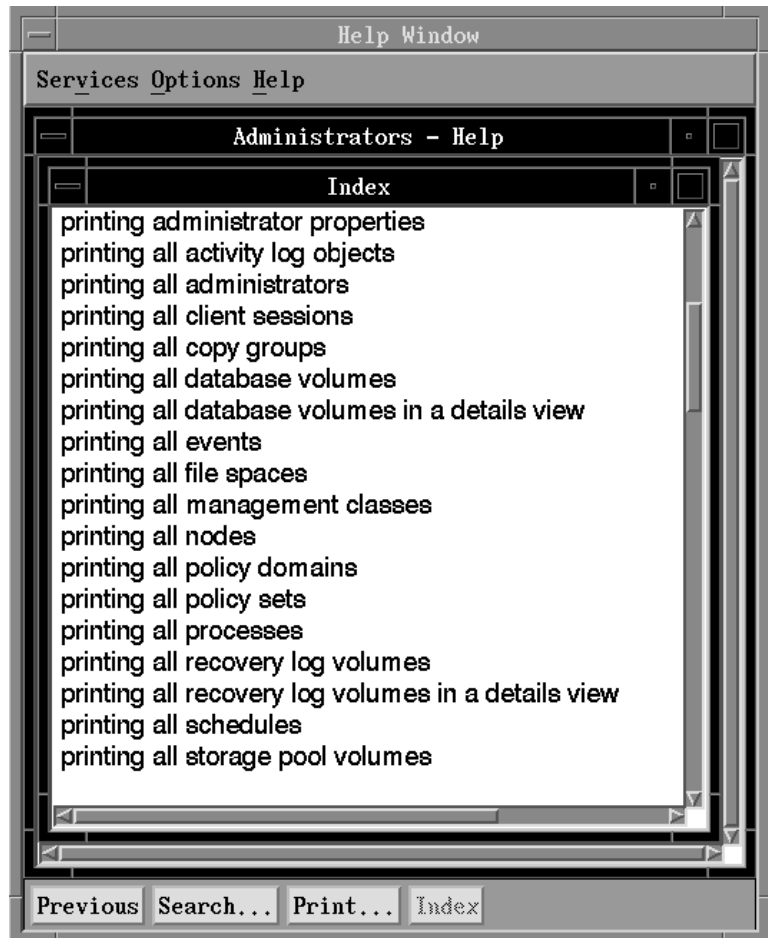
**4**

Type in the search word or words (for example, **print**).

**5**

Select the **Search** push button.

The search results appear in a help window.





You also can search for every occurrence of a word or phrase in one or more help topics by using the radio buttons in the Search window.

For example, to search the help index for information about printing:

### 1

Type **print** in the Search string field.

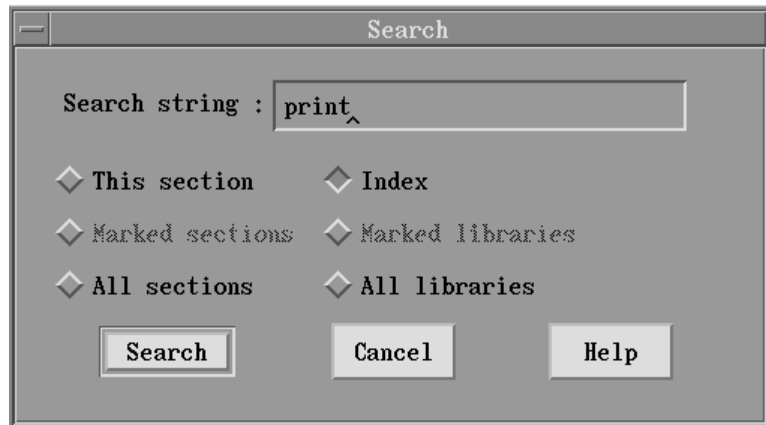
### 2

Select the **Index** radio button.

### 3

Select the **Search** push button.

The search results appear in a new help window.



To indicate to ADSM where to look for a word or phrase:

- Choose the **This section** radio button for the help topic that appears in the active help window.
- Choose the **Marked sections** radio button for selected help topics in the administrative graphical user interface.
- Choose the **All sections** radio button for all help topics in the administrative graphical user interface.
- Choose the **Index** radio button for the text of the index entries in the graphical user interface help.
- Choose the **Marked libraries** radio button for help topics in selected installed programs.
- Choose the **All libraries** radio button for help topics in all installed programs.

---

## Using the Administrative Command Line-Interface

The administrative command-line client lets administrators control and monitor the server through administrative commands. After you have installed the administrative client and modified the options file, registered administrators can access the server from any administrative client in the network. You can start an administrative client session in one of four modes:

- Console mode
- Mount mode
- Batch mode
- Interactive mode

You can also issue administrative commands from a server console session.

All these ways of accessing the server are discussed in this section. In addition, this section includes information about using administrative client options and macros.

For more information on using the command-line interface, see *ADSM Administrator's Reference*.

## Using Console Mode

Use console mode to monitor server activities as they occur. For example, you can monitor clients logging on to ADSM. This information is displayed on your terminal and, optionally, can be written to a file.

**Note:** You cannot enter administrator commands in console mode.

To start the administrative client in console mode, enter:

```
dsmadm -consolemode
```

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

```
dsmadm -id=yourid -password=secret -consolemode
```

To end an administrative client session, use one of the following keyboard break sequences:

<b>Environment</b>	<b>Break Sequence</b>
UNIX	Ctrl+C
CMS	HX
DOS	Ctrl+C, Ctrl+Break
OS/2	Ctrl+C, Ctrl+Break
TSO	ATTN
Windows and Windows NT	Ctrl+C, Ctrl+Break

## Using Mount Mode

Use mount mode to monitor removable media mount messages. This information is displayed on your terminal and, optionally, can be written to a file.

**Note:** You cannot enter administrator commands in mount mode.

To start the administrative client in mount mode, enter:

```
dsmdmc -mountmode
```

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

```
dsmdmc -id=yourid -password=secret -mountmode
```

To end an administrative client session in mount mode, use the Ctrl+C keyboard break sequence.

## Using Batch Mode

Use batch mode to enter a single command, which can be a MACRO command. Your administrative client session automatically ends when the command or macro has executed.

For example, to start the administrative client in batch mode and issue the ENABLE command, enter:

```
dsmdmc -id=smith -password=secret enable
```

## Using Interactive Mode

Use interactive mode to enter a series of administrative commands. To start an administrative client session in interactive mode, a server session must be available.

To start the administrative client in interactive mode, enter:

```
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 also enter your user ID and password in the DSMADM command by using the ID and PASSWORD options. For example, enter:

```
dsmadm -id=yourid -password=secret
```

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

```
quit
```

## Using Administrative Client Options

In all administrative client modes, the administrative client options modify your administrative client session responses. For example, the `-quiet` option specifies that you do not want ADSM to write any standard output messages to your terminal. For details about other options, see *ADSM Administrator's Reference*.

## Using Macros to Issue Commands

A macro is a file that contains one or more ADSM administrative commands. You can use macros when you want to issue commands repeatedly. You create a macro and issue the MACRO command with the name of the macro and, optionally, values for any substitution variables in the macro. For example, to issue the commands in the macro file named REGENG.MAC, enter:

```
macro regeng.mac
```

You can only issue macros from the administrative client in batch or interactive mode.

You can control whether the changes made to the database are permanent or not by using the COMMIT and ROLLBACK commands.

**COMMIT** Use this command to commit changes to the database made by the commands in a macro. For example:

```
/* macro to register policy administrators & grant authority*/  
register admin jones boat  
grant authority jones classes=policy  
commit  
register admin brown plane  
grant authority brown classes=policy  
commit
```

**ROLLBACK** Use this command to undo any changes to the database made by the commands in a macro, but not yet committed. By using the ROLLBACK command, you can test a macro without making the changes called for in the macro. For example:

```
/* macro to register policy administrators & grant authority*/
register admin jones boat
grant authority jones classes=policy
register admin brown plane
grant authority brown classes=policy
rollback
```

**Note:** If an administrative client session is running with the ITEMCOMMIT administrative client option, the ROLLBACK command has no effect.

For details about the MACRO command and the ITEMCOMMIT option, see *ADSM Administrator's Reference*.

## Using the Server Console Session

The *server console* is the terminal from which an administrator activates the server after ADSM is installed. ADSM has set up a special administrative user ID named SERVER\_CONSOLE that allows you to activate the server from the server console after ADSM is installed. SERVER\_CONSOLE is automatically registered as an administrator and is given system authority. Use the SERVER\_CONSOLE administrator user ID to register and grant system privileges to any other administrator as soon as ADSM is installed. (Authentication must be turned off to use SERVER\_CONSOLE as an administrative client ID.)

After you register other system administrators, reduce the authority of SERVER\_CONSOLE to operator privilege to restrict access to administrative functions. Run the server from one of your newly registered administrative clients, rather than continuing to use SERVER\_CONSOLE. For information on using the SERVER\_CONSOLE administrator user ID, see *ADSM Installing the Server and Administrative Client*.

You cannot modify SERVER\_CONSOLE attributes and therefore cannot make the following changes to this user ID:

- Register or update SERVER\_CONSOLE
- Lock or unlock SERVER\_CONSOLE from ADSM
- Rename SERVER\_CONSOLE
- Remove SERVER\_CONSOLE

The SERVER\_CONSOLE administrative ID does *not* receive a confirmation message when issuing commands that affect the availability of the server or data managed by the server. For more information on how to issue commands from the SERVER\_CONSOLE administrative ID, refer to *ADSM Administrator's Reference*.

You should change the privilege class of the SERVER\_CONSOLE administrative ID to operator privilege in order to prevent your tape operators from having system access to the server.

---

## Application Programming Interface

ADSM provides an application programming interface (API) that can be used by software vendors to integrate a storage management solution with existing applications to ensure that critical data is protected and easily recoverable.

When an application uses the ADSM API, it becomes an ADSM application client that can communicate with an ADSM server to backup, archive, or recover objects from ADSM storage.

An example of an application client is IBM Database 2\* AIX (DB2/6000\*) for the RISC System/6000, which provides online backup services through its own interface.

After an application client is installed, a user must modify the client options file to identify the node name of the workstation and the communication method used to communicate with the server.

Applications residing on AIX, HP-UX, OS/2, SunOS/Solaris, or Windows can be developed or modified to use the ADSM Storage Management API.

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 programming 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.

---

## Task, GUI, and Command Cross-Reference

The following tables list administrator tasks and show where they can be performed on the graphical user interface and command-line interface. The sequence of windows is shown under “Location in the GUI.” You can find detailed help for using the graphical user interface in its online help facility. For details about the commands, refer to the *ADSM Administrator’s Reference*, and the command-line interface online help (accessed through the HELP command).

## Tasks Found in “Using ADSTAR Distributed Storage Manager Interfaces”

Table 1 shows a listing of tasks and commands referenced in Chapter 2, “Using ADSTAR Distributed Storage Manager Interfaces” on page 7.

<i>Table 1. Tasks Found in “Using ADSTAR Distributed Storage Manager Interfaces”</i>		
<b>Task</b>	<b>Location in the GUI</b>	<b>Command Used</b>
Start the administrative client. (See “Using the Administrative Command Line-Interface” on page 22)	Not available	DSMADMC
End the interactive mode of the administrative command-line client. (See “Using Interactive Mode” on page 23.)	Not available	QUIT
Invoke a macro file of one or more ADSM commands. (See “Using Macros to Issue Commands” on page 24.)	Not available	MACRO
Commit changes made by commands. (See “Using Macros to Issue Commands” on page 24.)	Automatic	COMMIT
Rollback uncommitted changes made by commands. (See “Using Macros to Issue Commands” on page 24.)	Not available	ROLLBACK



## Tasks Found in “Managing Server Operations”

Table 2 shows a listing of tasks and commands referenced in Chapter 3, “Managing Server Operations” on page 49.

<i>Table 2 (Page 1 of 2). Tasks Found in “Managing Server Operations”</i>		
<b>Task</b>	<b>Location in the GUI</b>	<b>Command Used</b>
Start or restart the ADSM server. (See “Starting the Server” on page 49)	Not available	DSMSERV
Shut down the server. (See “Halting the Server” on page 50.)	Server	HALT
Query one or more client sessions. (See “Requesting Information about Client Sessions” on page 51.)	1. Server 2. Sessions	QUERY SESSION
Cancel one or more client sessions. (See “Canceling a Client Session” on page 52.)	1. Server 2. Sessions	CANCEL SESSION
Temporarily prevent client node access to the server. (See “Disabling or Enabling Server Access” on page 53.)	Server	DISABLE
Resume user activity on the server. (See “Disabling or Enabling Server Access” on page 53.)	Server	ENABLE
Request information about server background processes. (See “Requesting Information about Server Processes” on page 54.)	1. Server 2. Processes	QUERY PROCESS
Cancel a server background processes. (See “Canceling Server Processes” on page 54.)	1. Server 2. Processes	CANCEL PROCESS
Bring a random access volume online or offline. (See “Varying Disk Volumes Online or Offline” on page 55.)	<i>Database volumes:</i> 1. Database 2. Database Volumes <i>Recovery log volumes:</i> 1. Database Recovery Log 2. Recovery Log Volumes <i>Storage pool volumes:</i> 1. Storage Pools 2. Storage Pool Volumes	VARY

Table 2 (Page 2 of 2). Tasks Found in "Managing Server Operations"

<b>Task</b>	<b>Location in the GUI</b>	<b>Command Used</b>
Query system parameters. (See "Requesting Information about Server Status" on page 55.)	Server	QUERY STATUS
Specify the server name. (See "Setting the Server Name" on page 56)	Server	SET SERVERNAME
Query one or more server options. (See "Querying Server Options" on page 56.)	Not available	QUERY OPTION
Set the retention period for the activity log. (See "Setting the Activity Log Retention Period" on page 57.)	Server	SET ACTLOGRETENTION
Search activity log for messages. (See "Requesting Information from the Activity Log" on page 58.)	1. Server 2. Activity Log	QUERY ACTLOG
Set accounting records on or off. (See "Monitoring Accounting Records" on page 59.)	Server	SET ACCOUNTING
Get help on commands and error messages. (See "Getting Help on Commands and Error Messages" on page 61.)	Not available	HELP

## Tasks Found in “Managing the Database and Recovery Log”

Table 3 shows a listing of tasks and commands referenced in Chapter 4, “Managing the Database and Recovery Log” on page 63.

<i>Table 3 (Page 1 of 2). Tasks Found in “Managing the Database and Recovery Log”</i>		
<b>Task</b>	<b>Location in the GUI</b>	<b>Command Used</b>
Reset the utilization statistics for the database. (See “Monitoring Utilization” on page 70.)	Database	RESET DBMAXUTILIZATION
Reset the utilization statistics for the log. (See “Monitoring Utilization” on page 70.)	Database Recovery Log	RESET LOGMAXUTILIZATION
Define a database volume. (See “Step 1: Defining Disk Volumes” on page 75.)	1. Database 2. Database Volumes	DEFINE DBVOLUME
Define a recovery log volume. (See “Step 1: Defining Disk Volumes” on page 75.)	1. Database Recovery Log 2. Recovery Log Volumes	DEFINE LOGVOLUME
Increase the assigned capacity of the database. (See “Step 3: Extending the Capacity of the Database or Recovery Log” on page 77.)	Database	EXTEND DB
Increase the assigned capacity of the recovery log. (See “Step 3: Extending the Capacity of the Database or Recovery Log” on page 77.)	Database Recovery Log	EXTEND LOG
Display information about volumes defined to the database. (See “Step 1: Determining the Size of Volumes” on page 79.)	1. Database 2. Database Volumes	QUERY DBVOLUME
Display information about volumes defined to the recovery log. (See “Step 1: Determining the Size of Volumes” on page 79.)	1. Database Recovery Log 2. Recovery Log Volumes	QUERY LOGVOLUME
Decrease the assigned capacity of the database. (See “Step 3: Reducing the Capacity of the Database or Recovery Log” on page 82.)	Database	REDUCE DB

Table 3 (Page 2 of 2). Tasks Found in “Managing the Database and Recovery Log”

<b>Task</b>	<b>Location in the GUI</b>	<b>Command Used</b>
Decrease the assigned capacity of the recovery log. (See “Step 3: Reducing the Capacity of the Database or Recovery Log” on page 82.)	Database Recovery Log	REDUCE LOG
Delete a database volume. (See “Step 4: Deleting a Volume from the Database or Recovery Log” on page 82.)	1. Database 2. Database Volumes	DELETE DBVOLUME
Delete a recovery log volume. (See “Step 4: Deleting a Volume from the Database or Recovery Log” on page 82.)	1. Database Recovery Log 2. Recovery Log Volumes	DELETE LOGVOLUME
Display information on the database. (See “Requesting Information about the Database Buffer Pool” on page 84.)	Database	QUERY DB
Reset the buffer pool statistics for the database. (See “Resetting Database Buffer Pool Statistics” on page 85.)	Not available	RESET BUFPOOL
Display information on the recovery log. (See “Requesting Information about the Recovery Log Buffer Pool” on page 86.)	Database Recovery Log	QUERY LOG

## Tasks Found in “Managing Licensing, Privilege Classes, and Registration”

Table 4 shows a listing of tasks and commands referenced in Chapter 5, “Managing Licensing, Privilege Classes, and Registration” on page 89.

<i>Table 4 (Page 1 of 3). Tasks Found in “Managing Licensing, Privilege Classes, and Registration”</i>		
<b>Task</b>	<b>Location in the GUI</b>	<b>Command Used</b>
Register a new license. (See “Managing ADSM Licenses” on page 89.)	Server	REGISTER LICENSE
Audit the current server configuration and licenses. (See “Monitoring Licenses” on page 92.)	Server	AUDIT LICENSES
Display license information. (See “Monitoring Licenses” on page 92.)	Server	QUERY LICENSE
Set license audit period. (See “Monitoring Licenses” on page 92.)	Server	SET LICENSEAUDITPERIOD
Set password expiration date. (See “Setting User Password Expiration” on page 93.)	Server	SET PASSEXP
Set password authentication. (See “Setting Client Password Authentication” on page 93.)	Server	SET AUTHENTICATION
Register an administrator. (See “Registering Administrators or Updating Information” on page 93.)	Administrators	REGISTER ADMIN
Update an administrator. (See “Registering Administrators or Updating Information” on page 93.)	Administrators	UPDATE ADMIN
Add administrator authority. (See “Granting Administrative Authority” on page 94.)	Administrators	GRANT AUTHORITY
Revoke or reduce administrator authority. (See “Revoking or Reducing Administrative Authority” on page 98.)	Administrators	REVOKE AUTHORITY
Rename an administrator. (See “Renaming an Administrator” on page 101.)	Administrators	RENAME ADMIN

Table 4 (Page 2 of 3). Tasks Found in “Managing Licensing, Privilege Classes, and Registration”

<b>Task</b>	<b>Location in the GUI</b>	<b>Command Used</b>
Delete an administrator. (See “Removing Administrators” on page 101.)	Administrators	REMOVE ADMIN
Lock out an administrator. (See “Locking and Unlocking Administrators from the Server” on page 101.)	Administrators	LOCK ADMIN
Unlock an administrator. (See “Locking and Unlocking Administrators from the Server” on page 101.)	Administrators	UNLOCK ADMIN
Display information on one or more administrators. (See “Requesting Information about Administrators” on page 102.)	Administrators	QUERY ADMIN
Set open or closed registration. (See “Setting Client Node Registration” on page 103.)	Central Scheduler	SET REGISTRATION
Register a client node. (See “Registering Client Nodes” on page 104.)	Nodes	REGISTER NODE
Update a client node. (See “Updating Client Node Information” on page 105.)	Nodes	UPDATE NODE
Rename a client node. (See “Renaming Client Nodes” on page 106.)	Nodes	RENAME NODE
Lock out a client node. (See “Locking and Unlocking Client Nodes” on page 106.)	Nodes	LOCK NODE
Unlock a client node. (See “Locking and Unlocking Client Nodes” on page 106.)	Nodes	UNLOCK NODE
Query one or more client nodes. (See “Requesting Information about Client Nodes” on page 107.)	Nodes	QUERY NODE
Query one or more file spaces. (See “Requesting File Space Information” on page 108.)	File Spaces	QUERY FILESPACE
Delete client node data from the server. (See “Deleting a File Space” on page 109.)	File Spaces	DELETE FILESPACE

Table 4 (Page 3 of 3). Tasks Found in "Managing Licensing, Privilege Classes, and Registration"

<b>Task</b>	<b>Location in the GUI</b>	<b>Command Used</b>
Delete a client node. (See "Removing Client Nodes" on page 110.)	Nodes	REMOVE NODE
Register an application programming interface. (See "Registering an Application Programming Interface to the Server" on page 110.)	Nodes	REGISTER NODE

## Tasks Found in “Managing Policies”

Table 5 shows a listing of tasks and commands referenced in Chapter 6, “Managing Policies” on page 113.

<i>Table 5 (Page 1 of 3). Tasks Found in “Managing Policies”</i>		
<b>Task</b>	<b>Location in the GUI</b>	<b>Command Used</b>
Define a new policy domain. (See “Example: Defining a Policy Domain” on page 126.)	Policy Domains	DEFINE DOMAIN
Copy a policy domain. (See “Example: Defining a Policy Domain” on page 126.)	Policy Domains	COPY DOMAIN
Update a policy domain. (See “Example: Defining a Policy Domain” on page 126.)	Policy Domains	UPDATE DOMAIN
Copy a policy set. (See “Defining and Updating a Policy Set” on page 127.)	1. Policy Domains 2. Policy Sets	COPY POLICYSET
Define a policy set. (See “Defining and Updating a Policy Set” on page 127.)	1. Policy Domains 2. Policy Sets	DEFINE POLICYSET
Update a policy set. (See “Defining and Updating a Policy Set” on page 127.)	1. Policy Domains 2. Policy Sets	UPDATE POLICYSET
Define a management class. (See “Defining and Updating a Management Class” on page 127.)	1. Policy Domains 2. Management Class	DEFINE MGMTCLASS
Copy a management class. (See “Example: Define a New Management Class” on page 128.)	1. Policy Domains 2. Management Class	COPY MGMTCLASS
Update a management class (See “Example: Define a New Management Class” on page 128.)	1. Policy Domains 2. Management Class	UPDATE MGMTCLASS
Define a backup copy group. (See “Defining and Updating a Backup Copy Group” on page 128.)	1. Policy Domains 2. Backup Copy Groups	DEFINE COPYGROUP
Update a backup copy group. (See “Defining and Updating a Backup Copy Group” on page 128.)	1. Policy Domains 2. Backup Copy Groups	UPDATE COPYGROUP



Table 5 (Page 2 of 3). Tasks Found in "Managing Policies"

<b>Task</b>	<b>Location in the GUI</b>	<b>Command Used</b>
Define an archive copy group. (See "Defining and Updating an Archive Copy Group" on page 132.)	1. Policy Domains 2. Archive Copy Groups	DEFINE COPYGROUP
Update an archive copy group. (See "Defining and Updating an Archive Copy Group" on page 132.)	1. Policy Domains 2. Archive Copy Groups	UPDATE COPYGROUP
Assign a default management class. (See "Assigning a Default Management Class" on page 134.)	1. Policy Domains 2. Management Class	ASSIGN DEFMGMTCLASS
Activate policy set. (See "Validating and Activating Policy Sets" on page 134.)	1. Policy Domains 2. Policy Sets	ACTIVATE POLICYSET
Verifies a policy set. (See "Validating and Activating Policy Sets" on page 134.)	1. Policy Domains 2. Policy Sets	VALIDATE POLICYSET
Start inventory expiration processing. (See "Starting Expiration Processing" on page 135.)	Not available	EXPIRE INVENTORY
Display information about a copy group. (See "Querying Copy Groups" on page 136.)	1. Policy Domains 2. Backup Copy Groups or Archive Copy Groups	QUERY COPYGROUP
Query a management class. (See "Querying Management Classes" on page 137.)	1. Policy Domains 2. Management Class	QUERY MGMTCLASS
Query a policy set. (See "Querying Policy Sets" on page 137.)	1. Policy Domains 2. Policy Sets	QUERY POLICYSET
Query one or more policy domains. (See "Querying Policy Domains" on page 138.)	Policy Domains	QUERY DOMAIN
Delete a copy group. (See "Deleting Copy Groups" on page 138.)	1. Policy Domains 2. Backup Copy Groups or Archive Copy Groups	DELETE COPYGROUP
Delete a management class. (See "Deleting Management Classes" on page 139.)	1. Policy Domains 2. Management Class	DELETE MGMTCLASS
Delete a policy set. (See "Deleting Policy Sets" on page 139.)	1. Policy Domains 2. Policy Sets	DELETE POLICYSET

*Table 5 (Page 3 of 3). Tasks Found in "Managing Policies"*

<b>Task</b>	<b>Location in the GUI</b>	<b>Command Used</b>
Delete a policy domain. (See "Deleting Policy Domains" on page 139.)	Policy Domains	DELETE DOMAIN

## Tasks Found in “Scheduling Operations”

Table 6 shows a listing of tasks and commands referenced in Chapter 7, “Scheduling Operations” on page 141.

<i>Table 6 (Page 1 of 2). Tasks Found in “Scheduling Operations”</i>		
<b>Task</b>	<b>Location in the GUI</b>	<b>Command Used</b>
Select a central scheduling mode. (See “Setting the Scheduling Mode on the Server” on page 144.)	Central Scheduler	SET SCHEDMODES
Set maximum scheduled sessions. (See “Setting the Maximum Percentage of Sessions for Scheduled Operations” on page 145.)	Central Scheduler	SET MAXSCHEDESESSIONS
Randomly distribute scheduled start times. (See “Randomizing Schedule Start Times” on page 146.)	Central Scheduler	SET RANDOMIZE
Control how often client nodes contact the server to perform scheduled operations. (See “Setting How Often Clients Query the Server” on page 147.)	Central Scheduler	SET QUERYSCHEDPERIOD
Set number of times scheduler retries commands. (See “Setting the Number of Command Retry Attempts” on page 147.)	Central Scheduler	SET MAXCMDRETRIES
Set time between retry attempts. (See “Setting the Amount of Time between Retry Attempts” on page 148.)	Central Scheduler	SET RETRYPERIOD
Define a backup or archive schedule. (See “Defining or Updating Schedules” on page 149.)	1. Central Scheduler 2. Backup/Archive Schedules	DEFINE SCHEDULE
Update a schedule. (See “Defining or Updating Schedules” on page 149.)	1. Central Scheduler 2. Administrative Command Schedules or Backup/Archive Schedules	UPDATE SCHEDULE
Copy a schedule. (See “Copying Schedules” on page 152.)	1. Central Scheduler 2. Administrative Command Schedules or Backup/Archive Schedules	COPY SCHEDULE

Table 6 (Page 2 of 2). Tasks Found in "Scheduling Operations"

<b>Task</b>	<b>Location in the GUI</b>	<b>Command Used</b>
Query one or more schedules. (See "Querying Schedules" on page 153.)	1. Central Scheduler 2. Administrative Command Schedules or Backup/Archive Schedules	QUERY SCHEDULE
Delete one or more schedules. (See "Deleting Schedules" on page 153.)	1. Central Scheduler 2. Administrative Command Schedules or Backup/Archive Schedules	DELETE SCHEDULE
Associate client nodes with a schedule. (See "Associating Client Nodes with Schedules" on page 153.)	1. Central Scheduler 2. Administrative Events or Backup/Archive Events	DEFINE ASSOCIATION
Query client node associations with a schedule. (See "Querying Associations" on page 154.)	1. Central Scheduler 2. Administrative Events or Backup/Archive Events	QUERY ASSOCIATION
Delete node associations with a schedule. (See "Deleting Associations" on page 154.)	1. Central Scheduler 2. Administrative Events or Backup/Archive Events	DELETE ASSOCIATION
Query scheduled and completed events. (See "Querying Event Records" on page 155.)	1. Central Scheduler. 2. Administrative Events or Backup/Archive Events	QUERY EVENT
Set the retention period for event records. (See "Setting the Event Record Retention Period" on page 156.)	1. Central Scheduler 2. Administrative Events or Backup/Archive Events	SET EVENTRETENTION
Delete event records. (See "Deleting Event Records" on page 156.)	1. Central Scheduler 2. Administrative Events or Backup/Archive Events	DELETE EVENT

## Tasks Found in “Managing Storage Devices”

Table 7 shows a listing of tasks and commands referenced in Chapter 9, “Managing Storage Devices” on page 161.

<b>Task</b>	<b>Location in the GUI</b>	<b>Command Used</b>
Define a device class for tape. (See “Defining and Updating CARTRIDGE Device Classes” on page 163.)	Not available	DEFINE DEVCLASS
Define a device class for FILE. (See “Defining FILE Device Classes” on page 168.)	Not available	DEFINE DEVCLASS
Display information on one or more device classes. (See “Requesting Information about a Device Classes” on page 168.)	Not available	QUERY DEVCLASS
Change the attributes of a device class. (See “Defining and Updating CARTRIDGE Device Classes” on page 163.)	Not available	UPDATE DEVCLASS
Delete a device class. (See “Deleting Device Classes” on page 169.)	Not available	DELETE DEVCLASS

## Tasks Found in “Managing Storage Pools”

Table 8 shows a listing of tasks and commands referenced in Chapter 10, “Managing Storage Pools” on page 171.

<i>Table 8. Tasks Found in “Managing Storage Pools”</i>		
<b>Task</b>	<b>Location in the GUI</b>	<b>Command Used</b>
Define a storage pool. (See “Defining or Updating Storage Pools” on page 194.)	Storage Pools	DEFINE STGPOOL
Change the attributes of a storage pool. (See “Defining or Updating Storage Pools” on page 194.)	Storage Pools	UPDATE STGPOOL
Back up a primary storage pool. (See “Backing Up Storage Pools” on page 200.)	Storage Pools	BACKUP STGPOOL
Query one or more storage pools. (See “Monitoring the Use of Storage Pool Space” on page 202.)	Storage Pools	QUERY STGPOOL
Cancel a migration process. (See “Canceling the Migration Process” on page 206.)	1. Server 2. Processes	CANCEL PROCESS
Display file space information by storage pool. (See “Amount of Space Used by Client Node” on page 208.)	Not available	QUERY OCCUPANCY
Delete a storage pool. (See “Deleting a Storage Pool” on page 210.)	Storage Pools	DELETE STGPOOL
Recreate files in a primary storage pool. (See “Restoring Storage Pools” on page 211.)	Storage Pools	RESTORE STGPOOL

## Tasks Found in “Managing Storage Pool Volumes”

Table 9 shows a listing of tasks and commands referenced in Chapter 11, “Managing Storage Pool Volumes” on page 213.

<b>Task</b>	<b>Location in the GUI</b>	<b>Command Used</b>
Define a volume in a storage pool. (See “Storage Pool Volumes” on page 213.)	1. Storage Pools 2. Storage Pool Volumes	DEFINE VOLUME
Change user access to a storage pool volume. (See “Defining Storage Pool Volumes” on page 216.)	1. Storage Pools 2. Storage Pool Volumes	UPDATE VOLUME
Query one or more storage pool volumes. (See “Requesting General Information about Storage Pool Volumes” on page 218.)	1. Storage Pools 2. Storage Pool Volumes	QUERY VOLUME
Verify database information for a storage pool volume. (See “Auditing a Storage Pool Volume” on page 220.)	1. Storage Pools 2. Storage Pool Volumes	AUDIT VOLUME
Query the contents of a storage pool volume. (See “Viewing a Standard Report on the Contents of a Volume” on page 224.)	1. Storage Pools 2. Storage Pool Volumes	QUERY CONTENT
Move files on a storage pool volume. (See “Moving Files from One Volume to Another Volume” on page 225.)	1. Storage Pools 2. Storage Pool Volumes	MOVE DATA
Query one or more server processes. (See “Requesting Information about the Data Movement Process” on page 228.)	1. Server 2. Processes	QUERY PROCESS
Delete a storage pool volume. (See “Deleting an Empty Storage Pool Volume” on page 229.)	1. Storage Pools 2. Storage Pool Volumes	DELETE VOLUME
Recreate files in a primary storage pool volume. (See “Restoring Storage Pool Volumes” on page 230.)	1. Storage Pools 2. Storage Pool Volumes	RESTORE VOLUME

## Tasks Found in “Managing Tape Operations”

Table 10 shows a listing of tasks and commands referenced in Chapter 12, “Managing Tape Operations” on page 233.

<b>Task</b>	<b>Location in the GUI</b>	<b>Command Used</b>
Cancel one or more mount requests. (See “Maintaining Information on Tape Condition” on page 234.)	Not available	CANCEL REQUEST
Dismount a volume by its real device address. (See “Handling Tape Mount Requests” on page 235.)	Not available	DISMOUNT DEVICE
Dismount a volume by volume name. (See “Handling Tape Mount Requests” on page 235.)	Not available	DISMOUNT VOLUME
Display information on mounted sequential access volumes. (See “Handling Tape Mount Requests” on page 235.)	Not available	QUERY MOUNT



## Tasks Found in “Exporting and Importing Data”

Table 11 shows a listing of tasks and commands referenced in Chapter 13, “Exporting and Importing Data” on page 241.

<b>Task</b>	<b>Location in the GUI</b>	<b>Command Used</b>
Copy server information to sequential media. (See “Preparing to Export or Import Data” on page 242.)	Server	EXPORT SERVER
Request information about an export or import process (See “Requesting Information about an Export or Import Process” on page 244.)	1. Server 2. Processes	QUERY PROCESS
Copy administrator information to sequential media. (See “Exporting Administrator Information” on page 249.)	Administrators	EXPORT ADMIN
Copy client node information to sequential media. (See “Exporting Client Node Information” on page 250.)	Nodes	EXPORT NODE
Copy policy information to sequential media. (See “Exporting Policy Information” on page 250.)	Policy Domains	EXPORT POLICY
Import administrator information. (See “Importing Data from Sequential Media Volumes” on page 251.)	Administrators	IMPORT ADMIN
Import client node information. (See “Importing Data from Sequential Media Volumes” on page 251.)	Nodes	IMPORT NODE
Import policy information. (See “Importing Data from Sequential Media Volumes” on page 251.)	Policy Domains	IMPORT POLICY
Import the server. (See “Importing Data from Sequential Media Volumes” on page 251.)	Server	IMPORT SERVER
Rename a client filespace on the server. (See “Renaming a File Space” on page 260.)	File spaces	RENAME FILESPACE

## Tasks Found in “Recovering Data”

Table 12 shows a listing of tasks and commands referenced in Chapter 14, “Recovering Data” on page 261.

<i>Table 12 (Page 1 of 3). Tasks Found in “Recovering Data”</i>		
<b>Task</b>	<b>Location in the GUI</b>	<b>Command Used</b>
Define a volume copy of a database volume. (See “Defining Database or Recovery Log Mirrored Volumes” on page 263.)	<ol style="list-style-type: none"> <li>1. Database</li> <li>2. Database Volumes</li> </ol>	DEFINE DBCOPY
Define a volume copy of a recovery log volume. (See “Defining Database or Recovery Log Mirrored Volumes” on page 263.)	<ol style="list-style-type: none"> <li>1. Database Recovery Log</li> <li>2. Recovery Log Volumes</li> </ol>	DEFINE LOGCOPY
Display information on one or more database volumes. (See “Requesting Information about Mirrored Volumes” on page 264.)	<ol style="list-style-type: none"> <li>1. Database.</li> <li>2. Database Volumes</li> </ol>	QUERY DBVOLUME
Display information on one or more log volumes. (See “Requesting Information about Mirrored Volumes” on page 264.)	<ol style="list-style-type: none"> <li>1. Database Recovery Log</li> <li>2. Recovery Log Volumes</li> </ol>	QUERY LOGVOLUME
Save sequential volume history information to one or more files. (See “Establishing Volume History Backup Files” on page 270.)	<ol style="list-style-type: none"> <li>1. Server</li> <li>2. Sequential Volume History</li> </ol>	BACKUP VOLHISTORY
Delete sequential volume history information collected by the server. (See “Establishing Volume History Backup Files” on page 270.)	<ol style="list-style-type: none"> <li>1. Server</li> <li>2. Sequential Volume History</li> </ol>	DELETE VOLHISTORY
Display sequential volume history information. (See “Establishing Volume History Backup Files” on page 270.)	<ol style="list-style-type: none"> <li>1. Server</li> <li>2. Sequential Volume History</li> </ol>	QUERY VOLHISTORY
Backup device configuration information. (See “Establishing Device Configuration Backup Files” on page 271.)	Server	BACKUP DEVCONFIG
Backup the database. (See “Performing the Database Backups” on page 273.)	Database	BACKUP DB

Table 12 (Page 2 of 3). Tasks Found in "Recovering Data"

<b>Task</b>	<b>Location in the GUI</b>	<b>Command Used</b>
Define automatic database backup. (See "Setting the Database Backup Trigger" on page 273.)	Database	DEFINE DBBACKUPTRIGGER
Delete automatic database backup. (See "Setting the Database Backup Trigger" on page 273.)	Database	DELETE DBBACKUPTRIGGER
Update the automatic database backup. (See "Setting the Database Backup Trigger" on page 273.)	Database	UPDATE DBBACKUPTRIGGER
Query the setting for the backup trigger. (See "Setting the Database Backup Trigger" on page 273.)	Database	QUERY DBBACKUPTRIGGER
Set the recovery log mode. (See "Setting the Recovery Log Mode" on page 275.)	Database Recovery Log	SET LOGMODE
Restore the database to a point in time. (See "Restoring a Database by Using Point-in-Time Recovery" on page 276.)	Not available	DSMSERV RESTORE DB
Backup a primary storage pool. (See "Using Storage Pool Backup Features" on page 280.)	Storage Pools	BACKUP STGPOOL
Restore a primary storage pool from a copy storage pool. (See "Recovering by Using Backed Up Copies of Storage Pools" on page 280.)	Storage Pools	RESTORE STGPOOL
Restore a volume from a copy storage pool. (See "Recovery from Media Loss" on page 284.)	1. Storage Pools 2. Storage Pool Volumes	RESTORE VOLUME
Do an offline dump of the database (See "Database Salvage Utilities" on page 284.)	Not available	DSMSERV DUMPDB
Initialize the database and recovery log (See "Database Salvage Utilities" on page 284.)	Not available	DSMSERV INSTALL
Reload a dumped database (See "Database Salvage Utilities" on page 284.)	Not available	DSMSERV LOADDB

Table 12 (Page 3 of 3). Tasks Found in "Recovering Data"

Task	Location in the GUI	Command Used
Verify that a reloaded database is returned to a consistent state (See "Database Salvage Utilities" on page 284.)	Not available	DSMSERV AUDITDB

---

## Chapter 3. Managing Server Operations

Administrators can manage server operations. These operations include such tasks as starting and halting the server, managing client sessions, and monitoring server information. The sections listed in the following table begins at the indicated pages.

Section	Page
<b>Tasks:</b>	
Starting, halting, or restarting the server	49
Managing client sessions	51
Disabling or enabling server access	53
Managing server processes	53
Varying disk volumes online or offline	55
Requesting information about server status	55
Setting the server name	56
Querying server options	56
Managing the activity log	56
Diagnosing problems and monitoring errors	58
Monitoring accounting records	59
Getting help on commands and error messages	61

Most tasks presented in this chapter can be performed by using either the graphical user interface or the command line interface. Table 2 on page 29 shows whether a task can be performed on the graphical user interface, the command line interface, or both.

For information about issuing ADSM commands, refer to *ADSM Administrator's Reference*. For help performing a task on the graphical user interface, see the procedure described in "Using Online Help" on page 15.

---

### Starting, Halting, and Restarting the Server

Task	Required Privilege Class
Start, halt, and restart the server	System or operator

#### Starting the Server

To start the server, edit and run the ANRSTART job from ADSM.SAMPLIB.

The parameters you must update before running this job are the following:

- The job card parameters
- Any DD statements necessary to reference libraries required by the server.

For example, the server requires a C/370 library. If your C/370 library is provided by the C/370 Library Version 2 product, then you must reference its run time libraries in the ANRSTART job stream.

The following figure is a sample ANRSTART job with the job statement parameters added and the DD statements added to reference the C/370 run time libraries.

---

```
//ADSM    JOB (999,POK), 'ADSM START', CLASS=A, MSGLEVEL=(1,1),
// MSGCLASS=T, TIME=1440
/*JOBPARM L=9999
//JOBLIB  DD DSN=ADSM.LINKLIB, DISP=SHR
//*****
//*
//* THE 2 FOLLOWING DATASETS ARE THE C370 V2 RUNTIME
//* LIBRARIES.
//*
//*****
//          DD DSN=EDC.V2R1M0.SEDCLINK, DISP=SHR
//          DD DSN=PLI.V2R3M0.SIBMLINK, DISP=SHR
//SERVER  EXEC PGM=DSMSERV, DYNAMNBR=300
//OPT     DD DSN=ADSM.DSMSERV.OPT, DISP=SHR
//DSMAMENG DD DSN=ADSM.ANRMSG(ANRMENU), DISP=SHR
//DSK     DD DSN=ADSM.DISKLOG, DISP=SHR
//HLPAMENG DD DUMMY DSN=ADSM.ANRHLP(ANRHENU), DISP=SHR
//SYSPRINT DD SYSOUT=*
```

---

This job starts the server as a batch job. When the startup completes, the system console displays the following message:

```
ANR0993I ADSM server initialization complete.
```

This message is followed by the ADSM prompt. The ADSM server is now up and running.

## Halting the Server

You can halt the server without warning if an unplanned operating system problem requires you to return control to the operating system.

When you halt the server, all processes are abruptly stopped and client sessions are canceled, even if they are not completed. When the server is halted, administrator activity is not possible.

If possible, halt the server only after current administrative and client node sessions have completed or canceled. To shut down the server without severely impacting administrative and client node activity with the server, you must:

1. Disable the server to prevent new client node sessions from starting, as described in “Disabling or Enabling Server Access” on page 53.
2. Query for session information to identify any existing administrative and client node sessions, as described in “Requesting Information about Client Sessions” on page 51.
3. Notify any existing administrative and client node sessions that you plan to shut down the server. ADSM does not provide a network notification facility; you must use external means to notify users.
4. Cancel any existing administrative or client node sessions, as described in “Canceling a Client Session” on page 52.
5. Halt the server to shut down all server operations by using the HALT command.

## Restarting the Server

To start the server after it has been halted, follow the instructions in “Starting the Server” on page 49.

When you restart the server after it has been halted, ADSM rolls back any operations that had been in process to ensure that the database remains in a consistent state.

---

## Managing Client Sessions

Task	Required Privilege Class
Display information about client sessions	Any administrator
Cancel a client session	System or operator

A *client session* can be either an administrative or a client node session.

If you want to prevent clients from accessing the server for an extended period of time, use the LOCK and UNLOCK commands for client node and administrator sessions or disable the server.

For information on locking or unlocking administrators from the server, see “Locking and Unlocking Administrators from the Server” on page 101. For information on locking or unlocking client nodes from the server, see “Locking and Unlocking Client Nodes” on page 106.

## Requesting Information about Client Sessions

When administrators or users log on to the server, an administrative or client node session is established with the server. Each client session is assigned a unique session number.

To request information about client sessions, enter:

```
query session
```

The following figure shows a sample client session report.

Sess Number	Comm. Method	Sess State	Wait Time	Bytes Sent	Bytes Recvd	Sess Type	Platform	Client Name
3	Tcp/Ip	IdleW	9 S	7.8 K	706	Admin	OS/2	TOMC
5	Tcp/Ip	IdleW	0 S	1.2 K	222	Admin	OS/2	GUEST
6	Tcp/Ip	Run	0 S	117	130	Admin	OS/2	MARIE

Figure 2. Information about Client Sessions

Check the *session state* and *wait time* to determine the session state of the server and how long (in seconds, minutes, or hours) the session has been in the current state. The server session state can be one of the following:

<b>Start</b>	Connecting with a client session.
<b>Run</b>	Executing a client request.
<b>End</b>	Ending a client session.
<b>RecvW</b>	Waiting to receive an expected message from the client while a database transaction is in progress. A session in this state is subject to the COMMTIMEOUT limit.
<b>SendW</b>	Waiting for acknowledgement that the client has received a message sent by the server.
<b>MediaW</b>	Waiting for removable media to become available.
<b>IdleW</b>	Waiting for communication from the client, and a database transaction is NOT in progress. A session in this state is subject to the IDLETIMEOUT limit.

If a client does not initiate communication within the specified time limit set by the IDLETIMEOUT option in the server options file, then ADSM cancels the client session.

For example, if the IDLETIMEOUT option is set to 30 minutes, and a user does not initiate any operations within those 30 minutes, then ADSM cancels the client session. The client session is automatically reconnected to the server when it starts to send data again.

## Canceling a Client Session

You may cancel a client session when:

- A user is unable to continue with work because the system is not responding
- You want all sessions cancelled before halting the server

To cancel a client session, you must identify it by session number. You can display a session number by issuing a query for session information. For example, if the session number is 6, you cancel that session by entering:

```
cancel session 6
```

If you want to cancel all backup and archive sessions, enter:

```
cancel session all
```

If an operation, such as a backup or an archive process, is interrupted when you cancel the session, ADSM rolls back the results of the current transaction. That is, any changes made by the operation that are not yet committed to the database are undone. If necessary, the cancellation process may be delayed.

When user and administrator sessions are cancelled, those persons must log on to the server again. If they were in the process of performing a function when the session was cancelled, they must reissue their last command.

If the session you cancel is currently waiting for a media mount, the mount request is automatically cancelled.



If the session is in the Run state when it is canceled, the cancellation process does not take place until the session enters the SendW, RecvW, or IdleW state.

---

## Disabling or Enabling Server Access

Task	Required Privilege Class
Disable and enable client node access to the server	System or operator
Display server status	Any administrator

Disabling the server prevents users from establishing client node sessions with the server and suspends migration and reclamation processes. To disable the server, enter:

```
disable
```

When you disable the server, administrators can still access it, and current client node activity completes unless the user logs off or you cancel the client node session.

After the server has been disabled, you can enable the server to resume normal operations and allow users to access it by entering:

```
enable
```

You can issue the QUERY STATUS command to determine if the server is enabled or disabled.

---

## Managing Server Processes

Task	Required Privilege Class
Display information about a server background process	Any administrator
Cancel a server process	System

When a user or administrator issues an ADSM command or uses a graphical user interface to perform an operation, the server initiates a process, such as registering a client node, deleting a management class, or canceling a client session.

Many processes occur quickly and are run in the foreground, while others take longer to complete. To allow you to perform other tasks during long-running operations, ADSM runs the following operations as background processes:

- Auditing licenses
- Auditing a volume
- Backing up the database
- Backing up a storage pool
- Defining a database copy
- Defining a recovery log copy

- Deleting a file space
- Deleting a database volume
- Deleting a recovery log volume
- Deleting a storage volume
- Expiring the inventory
- Exporting or importing data
- Extending the database or recovery log
- Migrating files from one storage pool to the next storage pool
- Moving data from a storage volume
- Reclaiming space from tape storage volumes
- Reducing the database or recovery log
- Restoring a storage pool
- Restoring a volume
- Varying a database or recovery log volume online

The server assigns each background process an ID number and displays the process ID when the operation starts. For example, if you issue an EXPORT NODE command, ADSM displays a message similar to the following:

```
EXPORT NODE started as Process 10
```

## Requesting Information about Server Processes

You can request information about server background processes. If you know the process ID number, you can use the number to limit the search. However, if you do not know the process ID, you can display information about all background processes by entering:

```
query process
```

The following figure shows a server background process report after a DELETE FILESPACE command was issued. The report displays a process ID number, a description and a completion status for each background process.

Process Number	Process Description	Status
2	DELETE FILESPACE	Deleting filespace DRIVE_D for node CLIENT1: 172 files deleted.

Figure 3. Information about Background Processes

## Canceling Server Processes

You can cancel a server background process by specifying its ID number in the following command:

```
cancel process 2
```

You can issue the QUERY PROCESS command to find the process number. See “Requesting Information about Server Processes” for details.

If the process you want to cancel is currently waiting for a tape volume to be mounted (for example, a process initiated by EXPORT, IMPORT, or MOVE DATA commands), the mount request is automatically cancelled.

---

## Varying Disk Volumes Online or Offline

Task	Required Privilege Class
Vary a disk volume online or offline	System or operator

To perform maintenance on a disk volume or to upgrade DASD, you can vary a disk volume offline. For example, to vary the disk volume named ADSM.STORAGE.POOL001 offline, enter:

```
vary offline adsm.storage.pool001
```

If ADSM encounters a problem with a disk volume, the server automatically varies the volume offline.

After you have replaced the disk volume, you can make it available to the server by varying the volume online. For example, to make the disk volume named ADSM.STORAGE.POOL001 available to the server, enter:

```
vary online adsm.storage.pool001
```

---

## Requesting Information about Server Status

Any administrator can request information about the general server parameters defined by SET commands. To query the status of the server, enter:

```
query status
```

ADSM displays information about the server, such as:

- When the server was installed
- Whether the server is enabled or disabled
- Whether client registration is open or closed
- Whether passwords are required for client/server authentication
- How long passwords are valid
- Whether accounting records are being generated
- How long messages remain in the activity log before being deleted
- How many client sessions can concurrently communicate with the server
- How many client node sessions are available for scheduled work
- What percentage of the scheduling startup window is randomized
- What scheduling mode is being used

- How frequently client nodes can poll for scheduled work
- How many times and how frequently a client node can retry a failed attempt to perform a scheduled operation
- How long event records are retained in the database

---

## Setting the Server Name

Task	Required Privilege Class
Specify the server name	System

At installation, the server name is set to ADSM. After installation, you can use the SET SERVERNAME command to change the server name. You can use the QUERY STATUS command to see the name of the server.

To specify the server name as WELLS\_DESIGN\_DEPT., for example, enter the following:

```
set servername wells_design_dept.
```

---

## Querying Server Options

Any administrator can issue the QUERY OPTION command. Use the QUERY OPTION command to display information about one or more general server options.

You can issue the QUERY OPTION command with no operands to display general information about all defined server options. You also can issue the QUERY OPTION command with a specific options name or pattern-matching expression to display information on one or more options in the server options file.

To display general information about all defined server options, enter:

```
query option
```

---

## Managing the Activity Log

Task	Required Privilege Class
Change the size of the activity log	System or unrestricted storage
Set the activity log retention period	System
Monitor the activity log	Any administrator

The activity log contains all messages normally sent to the server console during server operation. The only exceptions are responses to commands entered at the console, such as responses to QUERY commands. Examples of messages sent to the activity log include:

- When client sessions start or end

- When migration starts and ends
- When backup versions are expired
- What data is exported to tape
- When expiration processing is performed
- What export or import processing is performed

Any error messages sent to the server console are also stored in the activity log.

Use the following sections to adjust the size of the activity log, set an activity log retention period, and request information about the activity log.

## Changing the Size of the Activity Log

Because the activity log is stored in the database, the size of the activity log should be factored into the amount of space allocated for the database, allowing at least 1MB of additional space for the activity log.

The size of your activity log depends on how many messages are generated by daily processing operations and how long you want to retain those messages in the activity log. When retention time is increased, the amount of accumulated data also increases requiring additional database storage.

When there is not enough space in the database or recovery log for activity log records, ADSM stops recording and sends messages to the server console. If you increase the size of the database or recovery log, ADSM starts activity log recording again. For information about increasing the size of the database or recovery log, see “Step 3: Extending the Capacity of the Database or Recovery Log” on page 77.

If you do not have enough space in the database for the activity log, you can do one of the following:

- Allocate more space to the database
- Reduce the length of time that messages are retained in the activity log

## Setting the Activity Log Retention Period

You can specify how long activity log information is retained in the database by using the SET ACTLOGRETENTION command.

The server automatically deletes messages from the activity log after they have passed the specified age. At installation, the activity log retention period is set to one day. To change the retention period to 30 days, for example, enter:

```
set actlogretention 30
```

You can display the current retention period for the activity log by querying the server status.

## Requesting Information from the Activity Log

You can request information stored in the activity log. To minimize processing time when querying the activity log, you can:

- Specify a time period in which messages have been generated. The default for the QUERY ACTLOG command shows all activities that have occurred in the previous hour.
- Specify the message number of a specific message or set of messages.
- Specify a string expression to search for specific text in messages.

For example, to review messages generated on May 30 between 8 a.m. and 5 p.m., enter:

```
query actlog begindate=05/30/1994 enddate=05/30/1994 -  
begintime=08:00 endtime=05:00
```

To request information about messages related to the expiration of files from the data storage inventory, enter:

```
query actlog msgno=0813
```

See the *ADSM Messages* for message numbers.

To request information about messages generated from the IMPORT NODE command, enter:

```
query actlog search='import node'
```

---

## Diagnosing Problems and Monitoring Errors

---

### Diagnosis, Modification or Tuning Information

---

At installation, a system programmer provides space in a flat file on your operating system for the error log, which is used by IBM service to help you diagnose severe errors if you are unable to access the activity log.

While you may find some entries in the error log useful for diagnosing and solving problems, you must contact a service representative for information about the syntax or meaning of the error log contents. If necessary, a service representative may ask you to issue TRACE commands to help diagnosis and solve server problems.

To ensure that the error log has sufficient space, you should prune the error log, deleting the oldest error messages, on occasion.

---

End of Diagnosis, Modification or Tuning Information

---

General-use programming interface

---

## Monitoring Accounting Records

Task	Required Privilege Class
Set accounting records on or off	System

ADSM accounting records show the server resources used during a session. This information lets you track the storage used by a client node session. At installation, accounting is set off. You can set accounting on by entering:

```
set accounting on
```

When accounting is set on, the server creates a session resource usage accounting record whenever a client node session ends. CPU time information is not recommended for use in billing as it does not account for all CPU time used for the client. In addition to time not recorded for database activity during the session, CPU time for migration, expiration, reclamation, and administrator activities on behalf of a client are not recorded.

Accounting records are stored in an SMF file as SMF record type 42 subtype 14 records:

Offset		Length	Field Type	Contents
Decimal	Hex			
0	0	2	Binary	Record length (standard SMF record header)
2	2	2	Binary	Segment descriptor (standard SMF record header)
4	4	1	Binary	System level identifier (standard SMF record header)
5	5	1	Binary	Record type, 42 (standard SMF record header)
6	6	4	Binary	Record time (standard SMF record header)
10	A	4	Packed	Record date (standard SMF record header)
14	E	4	EBCDIC	System ID (standard SMF record header)
18	12	4	EBCDIC	Subsystem ID (standard SMF record header)
22	16	2	Binary	Record subtype, 14
24	18	2	Binary	Number of triplets, 2 (offset/length/number of entries)
26	1A	2	N/A	Reserved
28	1C	4	Binary	Offset to product section
32	20	2	Binary	Length of product section, 12
34	22	2	Binary	Number of product sections, 1
36	24	4	Binary	Offset to ADSM section
40	28	2	Binary	Length of ADSM section
42	2A	2	Binary	Number of ADSM sections, 1

Offset		Length	Field Type	Contents
Decimal	Hex			
44	2C	2	EBCDIC	Product level
46	2E	8	EBCDIC	Product name, <i>ADSM</i>
54	36	2	N/A	Reserved
56	38	30	EBCDIC	Node name of ADSM client. If the node name does not fit completely within 30 characters, the client node name is listed as:  first...last  where <i>first</i> is the first 17 characters of the node name and <i>last</i> is the last 10 characters of the node name.
86	56	14	EBCDIC	Date and time of accounting (yyyymmddhhmmss)
100	64	4	Binary	Duration of session, in seconds
104	68	4	Binary	Number of archive database objects inserted during session
108	6C	4	Binary	Number of backup database objects inserted during session
112	70	4	Binary	Amount of archived files, in kilobytes, sent by the client to the server
116	74	4	Binary	Amount of backed up files, in kilobytes, sent by the client to the server
120	78	4	Binary	Amount of data, in kilobytes, communicated between a client node and the server during this session
124	7C	8	EBCDIC	Client owner name (UNIX)
132	84	8	EBCDIC	Node type.
140	8C	8	EBCDIC	Communication method used for session
148	94	4	Binary	Number of archive database objects retrieved during session
152	98	4	Binary	Number of backup database objects retrieved during session
156	9C	4	Binary	Amount of space, in kilobytes, retrieved by archived objects
160	A0	4	Binary	Amount of space, in kilobytes, retrieved by backed up objects
164	A4	4	Binary	Amount of Idle Wait time, in seconds, during the session
168	A8	4	Binary	Amount of Communications Wait time, in seconds, during the session
172	AC	4	Binary	Amount of Media Wait time, in seconds, during the session



Offset		Length	Field Type	Contents
Decimal	Hex			
176	B0	4	Binary	Amount of CPU time, in seconds, used by the server for basic client activity. This amount includes CPU time to send or receive data from the client but does not include CPU time to place data on, or retrieve it from database storage.
180	B4	1	Binary	Authentication method used
181	B5	1	Binary	Normal termination indicator (Normal=X'01')
182	B6	2	N/A	Reserved

\_\_\_\_\_ End of General-use programming interface \_\_\_\_\_

## Getting Help on Commands and Error Messages

Any administrator can issue the HELP command to display information about administrative commands and server and client messages.

You can issue the HELP command with no operands to display a menu of help selections. You also can issue the HELP command with operands that specify help menu numbers, commands and subcommands, or error message numbers.

To display the help menu, enter:

```
help
```

To display help information on the REMOVE commands, enter:

```
help remove
```

To display help information on a specific error message, such as ANR5300I for example, enter:

```
help 5300
```



---

## Chapter 4. Managing the Database and Recovery Log

Task	Required Privilege Class
Manage disk volumes used by the database and recovery log	System or unrestricted storage
Display information about the database and recovery log	Any administrator

Administrators can manage disk volumes used by the database and recovery log. The sections listed in the following table begin at the indicated pages.

Section	Page
<b>Concepts:</b>	
Database and recovery log	63
<b>Tasks:</b>	
Allocating space for the database and recovery log	65
Adding space to the database or recovery log	75
Deleting space from the database or recovery log	78
Optimizing the performance of the database or recovery log	83

Most tasks presented in this chapter can be performed by using either the graphical user interface or the command-line interface. Table 3 on page 31 shows whether a task can be performed on the graphical user interface, the command-line interface, or both.

For information about issuing ADSM commands, refer to *ADSM Administrator's Reference*. For help performing a task on the graphical user interface, see the procedure described in "Using Online Help" on page 15.

---

### Database and Recovery Log

The database manages information about the location of client backup and archive files residing in storage pools, and it records information necessary for ongoing server operations.

The recovery log is used to maintain a consistent database image by recording changes made to the database as a transaction proceeds. A transaction is an exchange between a client and the server that accomplishes a particular action. For example, defining a management class or copy group is a transaction.

If a transaction completes successfully, database changes are committed, and permanent changes are made to the database. If a transaction does not complete successfully, database changes are undone so that the database changes made by the failing transaction are removed. Because the database and recovery log contain references to transactions, it is important that they are considered together as a single unit.

This section describes the functions of the database and recovery log in more detail, as well as how they interact with each other.

## Database Information

The ADSM server database is used to record information about the location of client backup and archive files. It contains information on:

- Client nodes
- Administrators
- Policies
- Schedules
- Server settings
- Storage volume location for backed up, archived, and client-migrated files

The database also records the information needed for the following server operations:

- Access control information for administrative clients
- Activity logs, which contain the messages generated by the server
- Data storage inventory, used to locate the files that reside in storage pools
- Event records that are generated by the processing of scheduled commands
- Information about registered client nodes
- Information on ADSM volumes
- Policies assigned to the client nodes
- Schedules and their associations with client nodes

## Recovery Log Information

The recovery log is a buffer file within the server that:

- Manages updates that are a result of transactions from clients
- Keeps track of all the updates that have not yet been written to the database
- Plays a key role in restoring the database to a consistent state when a system failure occurs
- Knows which transactions have been completed so that they are copied into the recovered database
- Ensures that incomplete, uncommitted transactions are deleted from the database during the recovery process

## Relationship between the Database and the Recovery Log

Information stored in the server database, recovery log, and storage pool volumes is tightly related. The database and recovery log contain cross-references to units of work called *transactions*. The result of most transactions is that information, files, or instructions are sent to storage pool volumes. It is, therefore, important that all log references to this information be consistent so that the information can be located when it is needed. Examples of transactions include:

- Archiving a client file
- Associating a schedule with a client node
- Backing up a client file
- Defining a management class or copy group
- Deleting a volume from a disk storage pool
- Registering an administrator or a client node

The following is an overview of how the database and recovery log work together when transactions occur:

- A client requests a transaction, for example, a backup for a file.
- A transaction log record is written to the recovery log that describes the action.
- The recovery log associates, or cross-references, the log record with the actual transaction. This process allows the transaction to be rolled back, if necessary, during a recovery procedure.
- The transaction is eventually committed to the database; this process may not happen immediately.
- The transaction record within the recovery log is deleted; it does not necessarily occur the instant the transaction is committed.

“Optimizing the Performance of the Database or Recovery Log” on page 83 has more information on the steps that occur when transactions are processed.

## Storage Pool Volume Locations

The database contains reference information that points to the location of each client file held in storage pool volumes, making it possible to access all client files. As client files are moved, deleted, or migrated across storage volumes, the database is automatically updated to reflect the new location of the files. If the database or recovery log is lost, the data within the storage pools is no longer accessible because the pointer to its location are lost.

## Transaction Processing

When a client sends a transaction to the server, reference information is sent to the recovery log. However, the changes are not immediately made to the database. They are collected and held in a buffer pool and later updated to the database. Because of this, the database and recovery log are not always consistent. For example, when migration takes place, the information is in the recovery log but not immediately written to the database.

This difference between the database and recovery log is a concern only if you have hardware problems or a corruption of the database or storage pool volumes. The recovery process, which is described in Chapter 14, “Recovering Data” on page 261, restores the database to a state in which the database and recovery log are consistent. This state usually coincides with the point of the last database backup. During the recovery process, the server scans the database, ignoring the information in the recovery log and in cache memory. Therefore, any transactions held in the recovery log, but not yet written to the database, are lost.

---

## Allocating Space for the Database and Recovery Log

During installation, a system programmer allocates space by defining disk volumes to be used by the database and recovery log. The *ADSM Installing the Server and Administrative Client* describes how to allocate the minimum storage space required for the database and recovery log. You can use this minimum storage space to bring up the server and test your client/server environment.

As you register client nodes to the server, consider increasing the size of the database to meet the needs of your installation. With ADSM, you can add or delete storage space to the database or recovery log while the server is running.

Use this section to help you determine how much space you need for the database and recovery log. If you decide to add storage space, use this information to help you prepare disk volumes for database and recovery log storage by completing the following tasks:

1. Estimate the amount of space required for the database
2. Estimate the amount of space required for the recovery log
3. Understand how space is managed by the server
4. Allocate disk space to the database and recovery log

After you have completed these planning tasks, add storage space to the database and recovery log as described in “Adding Space to the Database or Recovery Log” on page 75.

## Step 1: Estimating the Amount of the Database Space

The server uses an internal *database* that stores all information related to its operation. Information stored in the database includes:

- Administrator and client node registration
- Storage management policy definitions
- Backup and archive schedules and associations
- Device class, storage pool, and volume definitions
- Information about backed up, archived, or client-migrated files in storage pools

Remember that the database tracks information *about* backup versions and archive copies. The actual data associated with backed up or archived files is stored in data storage; for example, on disk or tape volumes.

The capacity required for the database is largely affected by the number of files that are backed up or archived, because information about each file is recorded in the database. To estimate the amount of space required for the database, consider how many client nodes are registered to the server and how many files each client node might back up or archive. For information on estimating the amount of data backed up or archived by workstations, see “Estimating Space Needs for Storage Pools” on page 192.

As a general guideline, the initial amount of space allocated to the database should be between 1% to 5% the total amount of space required for data storage. For example, if your installation requires 10GB of data storage, then the size of your database should be between 100MB to 500MB.

If you back up primary storage pools to copy storage pools, the database requires additional space. The database requires a small amount of overhead space and about 200 bytes for each file copy in a copy storage pool.

## Step 2: Estimating the Amount of Recovery Log Space

The server uses the *recovery log* to keep a record of all changes made to the database. When changes occur, the server updates the recovery log before it updates the database. If a system failure occurs, the server uses the information in the recovery log to roll back uncommitted transactions from the database, ensuring that information in the database is returned to a consistent state.

Because the server supports concurrent client sessions, numerous updates can occur at the same time. For example, whenever a file is backed up or archived, a

number of records are added to the recovery log. Consequently, the size of the recovery log is dependent on the number of concurrent client sessions communicating with the server, rather than the number of files in data storage. The number of concurrent sessions and the number of background processes executing on the server determine the number of transactions.

The number of concurrent client sessions cannot be greater than the number of maximum client sessions, as defined by the MAXSESSIONS option in the server options file, as described in the *ADSM Installing the Server and Administrative Client*.

Initially, begin with at least 12MB for the recovery log. Then monitor the utilization of the recovery log to determine whether you should increase or decrease the size of the recovery log.

**Note:** The recovery log must be large enough to store additional recovery log records. Operating in roll-forward mode significantly increases recovery log storage requirements. The extent of the increase is determined by the number of database transactions since the last database backup. In roll-forward mode, the recovery log tracks all transactions since the last database backup. See “Database Backup and Recovery” on page 265 for details.

### Step 3: Understanding How Space is Managed by the Server

ADSM tracks all volumes defined to either the database or the recovery log as one *logical volume*. To determine how much space is available in each logical volume, query the database or recovery log.

To request information about the database, enter:

```
query db
```

The server displays a report, similar to Figure 4.

Available Space (MB)	Assigned Capacity (MB)	Maximum Extension (MB)	Maximum Reduction (MB)	Page Size (bytes)	Total Pages	Used Pages	%Util	Max. %Util
196	96	100	92	4,096	24,576	86	0.3	0.3

Figure 4. Information about the Database

To request information about the recovery log, enter:

```
query log
```

The server displays a report, similar to Figure 5 on page 68.

Available Space (MB)	Assigned Capacity (MB)	Maximum Extension (MB)	Maximum Reduction (MB)	Page Size (bytes)	Total Pages	Used Pages	%Util	Max. %Util
12	12	0	8	4,096	3,072	68	2.2	2.2

Figure 5. Information about the Recovery Log

When you query the database or query the recovery log, ADSM displays the following space utilization information:

- Available space
- Assigned capacity
- Utilization
- Maximum percent utilized

Use this information to determine whether you need to allocate more disk space and define additional volumes to the database and recovery log.

### Available Space

*Available space* is the amount of space, in megabytes, that is available to the database and recovery log. You can use this available space to:

- Extend the capacity of the database or recovery log
- Provide sufficient free space before you try to delete a volume from the database or recovery log

The server determines available space by adding the amount of usable space from all volumes defined to the database or recovery log. ADSM calculates the amount of usable space by:

- Subtracting 1MB from each volume for overhead processing
- Dividing the remaining space into 4MB partitions. Any remaining megabytes result in unused space on the volume

The following examples illustrate how you can efficiently allocate disk space so that most of the space is available to the database and recovery log.

**An Example of Poor Use of Disk Space:** If you allocate four 20MB volumes for the recovery log, the server:

- Subtracts 1MB from each volume for overhead processing so that each volume has 19MB of available space
- Divides the remaining space (19MB) into 4MB partitions so that each volume contains four 4MB partitions

In this case, the server uses 16MB of the remaining 19MB from each volume and leaves 3MB of unused space on every volume.

In this example, only 64MB of the initial 80MB are available to the recovery log.




**An Example of Better Use of Disk Space:** If you allocate four 25MB volumes for the database, the server:

- Subtracts 1MB from each volume for overhead processing so that each volume has 24MB of available space
- Divides the remaining space (24MB) into 4MB partitions so that each volume contains six 4MB partitions

In this case, the server uses all 24MB of the remaining 24MB in each volume.

In this example, 96MB of the initial 100MB are available for use by the database.

Figure 6 shows that the available space for the database logical volume is 96MB.



---

Figure 6. An Example of Available Space and Assigned Capacity

### Assigned Capacity

*Assigned capacity* is the portion of available space that can be used to store database or recovery log information. During installation, the server automatically extends the database and recovery log to match the available space, as shown in Figure 6.

After installation, if you decide to add additional volumes, then you must extend the database or recovery log to increase the assigned capacity. For information on extending the capacity of the database or recovery log, see “Step 3: Extending the Capacity of the Database or Recovery Log” on page 77.

### Utilization

As data is added to the database or recovery log, the server tracks the percentage of utilization. *Utilization* is the percent of assigned capacity used by the database or recovery log at a specific point in time.

During the day, the utilization percentage varies. The *maximum percent utilized* is the highest percent of assigned capacity used by the database or recovery log.

For example, Figure 7 on page 70 shows that 80% of a recovery log is utilized during the most active period of the day.




Figure 7. An Example of Utilization

## Monitoring Utilization

The maximum amount of space used by the recovery log can vary significantly throughout the day because it is proportional to the transaction load on the system. For example, if many backups occur in the morning, the amount of space required by the recovery log increases during this period.

The maximum amount of space used by the database is more consistent with the utilization percentage because the amount of database space consumed is proportional to the number of objects inserted into or deleted from the database. The utilization percentage and maximum percent utilized tend to be the same, unless a large number of objects are deleted from the database.

Database and recovery log utilization percentages are reset automatically whenever the server is restarted. You can also adjust the size of your database or recovery log by using the `RESET DBMAXUTILIZATION` and `RESET LOGMAXUTILIZATION` commands to meet your storage needs.

To monitor the utilization of the database or recovery log on a continuous basis, you can reset the maximum utilization counters. In this way, you can monitor how much space is being utilized every day.

For example, you might initially want to reset database and recovery log utilization statistics on a daily basis to set the maximum utilization percentage equal to the current utilization.

To reset the maximum utilization statistic for the database, enter:

```
reset dbmaxutilization
```

To reset the maximum utilization statistic for the recovery log, enter:

```
reset logmaxutilization
```

## Step 4: Allocating Space for the Database and Recovery Log

Before you allocate disk space for the database or recovery log, consider:

- How the placement of volumes affects the availability of the database and recovery log
- How the placement of volumes affects the performance of the database and recovery log
- How the size of volumes affects space utilization

Then allocate space on disk volumes for use by the database and recovery log.

### Ensuring the Availability of the Database and Recovery Log

To protect database and recovery log volumes from media failure, you must use the mirroring feature to protect the server from media failure. In addition, you want to place database and recovery log mirror volumes on separate physical disks. If you cannot assign each volume to its own HDA, assign database and recovery log volumes to separate physical disks. See “Database and Recovery Log Mirroring” on page 262 for information on the mirroring feature.

### Improving the Performance of the Database and Recovery Log

To improve performance, define more than one volume for the database and recovery log, and place these volumes on separate disks to allow simultaneous access to different parts of the database or recovery log.

When possible, assign each database and recovery log volume to separate DASD strings, preferably on DASD with different control units and channels.

### Using Space Efficiently

To use disk space efficiently, allocate a few large disk volumes rather than many small disk volumes. You want to allocate a few large disk volumes to avoid losing space to ADSM overhead processing. See “Available Space” on page 68 for information on allocating space more efficiently.

### Allocating Space on Disk Storage

If the SMS address space is not active, edit and run the ANRFMT job (see “Using the ANRFMT Job”) to allocate and pre-format the new data set. If the SMS address space is active, you can also use the ANRFMT2 job from the ADSM.SAMPLIB (see “Using the ANRFMT2 Job.” on page 73) to allocate and pre-format the new data set.

**Using the ANRFMT Job:** Edit the ANRFMT job from the ADSM.SAMPLIB as follows:

- Customize the JOB card parameters for your installation
- In the DEFINE CLUSTER statement:
  - Enter the name of the new data set in the NAME parameter
  - Enter the size in megabytes of the new data set in the LINEAR MB parameter
  - Enter the volume on which to allocate the new data set in the VOL parameter

- Enter the name of the new data set in the PARM parameter of the ADSMFMT EXEC statement

---

```

//ANRFMT JOB MSGLEVEL=(1,1),
//          MSGCLASS=H,CLASS=A,REGION=7M,TIME=1440
//*
//*****
//*
//* ANRFMT Use ANRFMT to extend the recovery log, database or a
//*       storage pool if SMS is not active on your system. If
//*       SMS is active, you might choose to use ANRFMT2 as a
//*       PROC.
//*
//* Dependency: none
//*
//* Function:
//* - allocate a VSAM linear dataset
//* - pre-format the dataset for use by the ADSM server
//*
//* Usage:
//* - customize the job card
//*
//* - modify:
//*   step      parameter value
//*   -----  -
//*   ALLOCAT   NAME      name of the VSAM dataset to allocate
//*                   and pre-format
//*                   MB      size of the dataset in megabytes
//*                   VOL     volume serial for the VSAM dataset
//*
//*   ADSMFMT   PARM=     dataset name here must match the NAME
//*                   field of the preceding allocation step
//*
//*****

```

---

Figure 8 (Part 1 of 2). ANRFMT Job from ADSM.SAMPLIB

---

```

//*****
//*
//* Allocate the VSAM data set
//*
//*****
//ALLOCAT EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=*
//SYSIN DD *
DEFINE CLUSTER( NAME(ADSM.DATASET.NAME) LINEAR MB(100) VOL(ADSM01) )
/*
//*
//*
//*****
//*
//* Pre-format the data set before defining it to the ADSM Server
//* Note: this step will not be executed if the ALLOCAT
//* step fails
//*
//*****
//ADSMFMT EXEC PGM=DSMFMT,DYNAMNBR=300,COND=(0,NE,ALLOCAT),
// PARM='/ADSM.DATASET.NAME'
//OPT DD DSN=ADSM.ANRSERV.OPTIONS,DISP=SHR
//DSMAMENG DD DSN=ADSM.ANRMSG(ANRMENU),DISP=SHR
//SYSPRINT DD SYSOUT=*
//SYSTEM DD SYSOUT=*
//

```

---

Figure 8 (Part 2 of 2). ANRFMT Job from ADSM.SAMPLIB

**Using the ANRFMT2 Job.:** If you install ANRFMT2 as a procedure on an appropriate library (for example, SYS1.PROCLIB), you can invoke the allocation and pre-formatting of a dataset by entering the following at the system console:

```
s anrfmt2,name=<name>,mb=<mb>,volser=<volser>
```

where:

- <name>** The new dataset to define for ADSM usage
- <mb>** The size in megabytes
- <volser>** The volume on which to allocate <name>

---

```

//ANRFMT2 PROC NAME=,MB=,VOLSER=
//*
//*****
//*
//* ANRFMT2 Use ANRFMT2 to extend the recovery log, database or a
//* storage pool.
//* ANRFMT2 is a proc: if installed in an appropriate
//* library, it may be started from the system console.
//*
//* Dependency: SMS must be active; the dataset does not have to
//* be system-managed
//*
//* Function:
//* - allocate a VSAM linear dataset
//* - pre-format the dataset for use by the ADSM server
//*
//* Usage:
//* - invoke via JCL (see ANRALL02 for examples)
//*
//* - OR, issue on the system console:
//* s anrfmt2,name=<name>,mb=<mb>,volser=<volser>
//*
//* where:
//* <name> is the new dataset to define for ADSM usage
//* <mb> is the size in megabytes
//* <volser> is the volume on which to allocate <name>
//*
//*****
//*****
//* allocate a new dataset
//*****
//ALLOCVOL EXEC PGM=IEFBR14
//DSMVOL DD DSN=&NAME,
// DISP=(NEW,CATLG),RECOG=LS,AVGREC=M,
// SPACE=(&MB,(1)),UNIT=SYSDA,VOL=SER=&VOLSER
//*
//*
//*****
//* pre-format the dataset before defining it to the ADSM server
//* Note: this step will not be executed if the ALLOCVOL
//* step fails
//*****
//DSMFMT EXEC PGM=DSMFMT,DYNAMNBR=300,COND=(0,NE,ALLOCVOL),
// TIME=1440,PARM='/&NAME'
//OPT DD DSN=ADSM.ANRSERV.OPTIONS,DISP=SHR
//DSMAMENG DD DSN=ADSM.ANRMSG(ANRMENU),DISP=SHR
//SYSPRINT DD SYSOUT=*
//SYSTEM DD SYSOUT=*
//

```

---

Figure 9. ANRFMT2 Job from ADSM.SAMPLIB

After you have allocated space to disk volumes, you can add them to the database or recovery log as described in “Adding Space to the Database or Recovery Log” on page 75.

---

## Adding Space to the Database or Recovery Log

**Warning:** The size of an allocated database, recovery log, or storage pool volume cannot be changed after it has been defined to the ADSM server. ADSM uses the initial size allocation of the volume at the time it is defined to the server to calculate data placement for later retrieval. If you change the size of ADSM volumes by using external operating system utilities, ADSM may not initialize correctly, and data could be lost.

To add space to the database or recovery log, you must first define a volume and then extend the capacity of the database or recovery log.

You can add space to the database or the recovery log while the server is running. Use this section to help you:

1. Define disk volumes to the database or recovery log
2. Determine the maximum extension of the database or recovery log
3. Extend the capacity of the database or recovery log

### Step 1: Defining Disk Volumes

During installation, you initially defined volumes for your database and recovery log. At any time you can increase the size of the database and recovery log by adding more volumes. Each time you define a volume to the database or recovery log, you increase the available space without interrupting server operations. You can use this space to:

- Extend the capacity of the database or recovery log to store additional database or recovery log records, as described in “Step 3: Extending the Capacity of the Database or Recovery Log” on page 77
- Provide sufficient space to delete a volume from the database or recovery log, as described in “Step 4: Deleting a Volume from the Database or Recovery Log” on page 82

Database volumes should be defined in increments of 4MB. If you define a 101MB disk volume name VOL5 to the database, the server tracks the volume as part of the database logical volume, as shown in Figure 10. Because 1MB from VOL5 is used for overhead process, 100MB is added to the database to increase the available space to 196MB. However, the assigned capacity remains at 96MB.




Figure 10. Adding Volumes Increases Available Space

To define additional database volumes, do the following steps:

1. Allocate and format the new volume by running one of the following two jobs:

If the SMS address space is *not* active, use the ANRFMT from the ADSM.SAMPLIB. For more information, see “Using the ANRFMT Job” on page 71.

If the SMS address space is active, use the ANRFMT2 if you have installed it in a procedure library. For more information on the ANRFMT2 job, see “Using the ANRFMT2 Job.” on page 73.

2. Define the volume. For example:

- To define the ADSM.DBVOL.JRM database volume enter:

```
DEFine DBVolume ADSM.DBVOL.JRM
```

- To define the ADSM.LOGVOL.JRM recovery log volume enter:

```
DEFine DBVolume ADSM.LOGVOL.JRM
```

3. To make the additional space available to the database, do the following:

- a. Determine the maximum extension size by entering the QUERY DB or QUERY LOG command.
- b. Extend the database or recovery log by 15MB by entering on of the following:

```
EXTEND DB 15  
  
or  
  
EXTEND LOG 15
```

## Step 2: Determining the Maximum Extension

After you define volumes to the database or recovery log, determine how much the database or recovery log can be extended.

The *maximum extension* is determined by subtracting the assigned capacity from the available space. For example, Figure 11 on page 77 shows that the maximum extension of the database is 100MB.



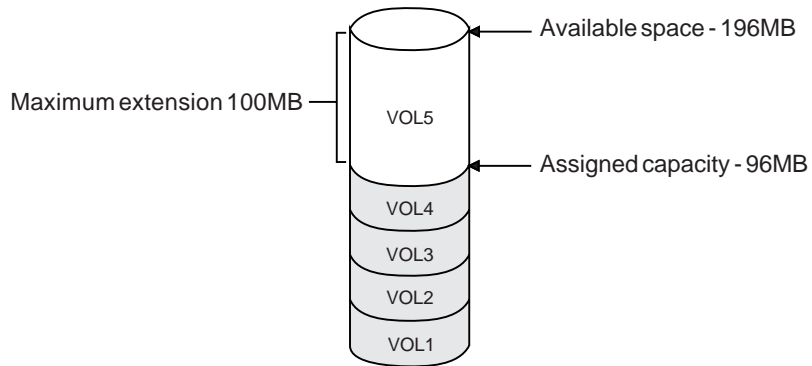


Figure 11. An Example of Maximum Extension

To determine the maximum extension of the database, enter:

```
query db
```

Figure 12 displays a standard database report, which shows that the available space is 196MB, the assigned capacity is 96MB, and the maximum extension is 100MB.

Available Space (MB)	Assigned Capacity (MB)	Maximum Extension (MB)	Maximum Reduction (MB)	Page Size (bytes)	Total Pages	Used Pages	%Util	Max. %Util
196	96	100	92	4,096	24,576	86	0.3	0.3

Figure 12. A Database Report to Determine Maximum Extension

### Step 3: Extending the Capacity of the Database or Recovery Log

To increase the capacity of database or recovery log, you must extend the database or recovery log in 4MB increments. If you do not specify the extension in 4MB increments, then ADSM rounds the number to the next higher 4MB partition. For example, if you extend the recovery log by 1MB, ADSM extends the capacity of the recovery log by 4MB.

For example, to increase the capacity of the database by 100MB, enter:

```
extend db 100
```

To increase the capacity of the recovery log by 4MB, enter:

```
extend log 4
```

When you extend the database or recovery log, ADSM starts a background process to format the new space. When the background process completes, the capacity of the database is increased by 100MB, as shown in Figure 13 on page 78.




Figure 13. Extending the Capacity of the Database

You can query the server to see the assigned capacity of the database or recovery log. For example, to view the assigned capacity of the database, enter:

```
query db
```

After you extend the database, the available space is 196MB, the assigned capacity is 196MB, and the maximum extension is 0MB, as shown in Figure 14.

Available Space (MB)	Assigned Capacity (MB)	Maximum Extension (MB)	Maximum Reduction (MB)	Page Size (bytes)	Total Pages	Used Pages	%Util	Max. %Util
196	196	0	192	4,096	50,176	111	0.2	0.2

Figure 14. A Database Report to Determine Assigned Capacity

After the capacity of the database or recovery log is extended, additional data can be added to the database or recovery log.

## Deleting Space from the Database or Recovery Log

You should delete volumes from the database or recovery log only if there is sufficient space on other volumes within the database or recovery log. To determine whether there is sufficient space to delete a volume, complete the following steps:

1. Determine the size of the volumes you want to delete.

Request a detailed database or recovery log volume report to view the available space, allocated space, and free space on each database or recovery log volume, as described in "Step 1: Determining the Size of Volumes" on page 79.

2. Determine whether the database or recovery log has sufficient *free space* on which to store information from the volume you want to delete, as described in “Step 2: Determining If There Is Sufficient Free Space” on page 81.
3. If you do not have sufficient free space, do one of the following:
  - Add more space to the database or recovery log by defining additional volumes, as described in “Step 1: Defining Disk Volumes” on page 75
  - Reduce the capacity of the database to free up existing space in the database or recovery log, as described in “Step 3: Reducing the Capacity of the Database or Recovery Log” on page 82
4. Delete volumes from the database or recovery log, when the amount of available space is equal or greater to the size of the volumes you want to delete, as described in “Step 4: Deleting a Volume from the Database or Recovery Log” on page 82.

## Step 1: Determining the Size of Volumes

Any administrator can request detailed information about volumes defined to the database or recovery log. Detailed reports are useful for determining how much space has been allocated and used by an individual volume in the database or recovery log.

For example, to display a detailed volume report about volumes defined to the database, enter:

```
query dbvolume format=detailed
```

Figure 15 on page 80 displays a detailed report about volumes defined to the recovery log. This example shows that VOL4, VOLD, and VOL300 are a group of mirrored volumes, and VOL5, VOLE, and VOL200 are a group of mirrored volumes.

This example also shows that there is no free space available on any of these volumes because all available space has been allocated to the database.

```

Volume Name (Copy 1): VOL4
      Copy Status: Sync'd
Volume Name (Copy 2): VOLD
      Copy Status: Sync'd
Volume Name (Copy 3): VOL300
      Copy Status: Sync'd
Available Space (MB): 24
Allocated Space (MB): 24
      Free Space (MB): 0

Volume Name (Copy 1): VOL5
      Copy Status: Sync'd
Volume Name (Copy 2): VOLE
      Copy Status: Sync'd
Volume Name (Copy 3): VOL200
      Copy Status: Sync'd
Available Space (MB): 100
Allocated Space (MB): 100
      Free Space (MB): 0

```

*Figure 15. Information about Database Volumes*

To display a detailed volume report about the log volume defined to the database, enter:

```
query logvolume format=detailed
```

Figure 16 displays a detailed report about the recovery log volume.

```

Volume Name (Copy 1): /home/bill/dsmserv/build/log.1
      Copy Status: Sync'd
Volume Name (Copy 2):
      Copy Status: Undefined
Volume Name (Copy 3):
      Copy Status: Undefined
Available Space (MB): 8
Allocated Space (MB): 8
      Free Space (MB): 0

Volume Name (Copy 1): VOL5
      Copy Status: Sync'd
Volume Name (Copy 2): VOLE
      Copy Status: Sync'd
Volume Name (Copy 3): VOL200
      Copy Status: Sync'd
Available Space (MB): 100
Allocated Space (MB): 100
      Free Space (MB): 0

```

*Figure 16. Information about Database Volumes*

## Step 2: Determining If There Is Sufficient Free Space

To determine if there is sufficient free space to delete a volume from the database or recovery log, request information about the database or recovery log.

For example, to determine how much free space is available in the database, enter:

```
query db
```

Figure 17 displays a standard database report.

Available Space (MB)	Assigned Capacity (MB)	Maximum Extension (MB)	Maximum Reduction (MB)	Page Size (bytes)	Total Pages	Used Pages	%Util	Max. %Util
196	196	0	176	4,096	50,176	4,755	9.5	9.5

Figure 17. Information about the Database

From this report, you can determine if there is sufficient free space in one of two ways:

- Subtract the assigned capacity from the available space.  
In this example,  $196\text{MB} - 196\text{MB} = 0$ .
- Determine how much assigned capacity is not being used to store data as shown by the maximum reduction column.

The server automatically calculates the space actually used and reports the amount of free space in the *Maximum Reduction* column.

In this example, the maximum reduction is 176MB, as shown in Figure 18.

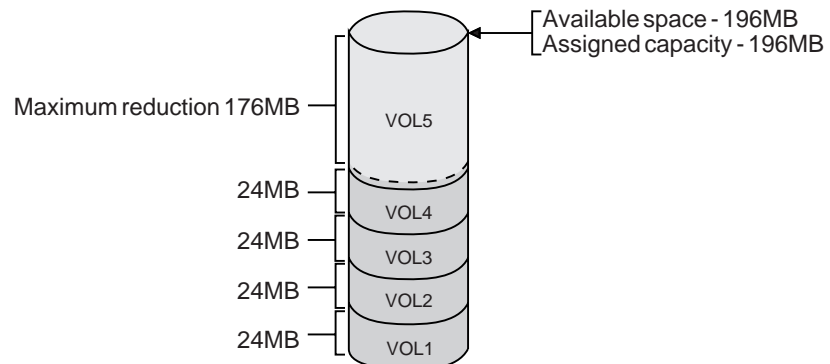


Figure 18. Understanding Maximum Reduction

### Step 3: Reducing the Capacity of the Database or Recovery Log

The *maximum reduction* identifies the number of megabytes by which you can reduce the database or recovery log. By reducing the database or recovery log, you might be able to free up enough space to allow you to delete a volume from the database or recovery log.

You can reduce the capacity of the database or recovery log in 4MB increments. If you do not reduce by increments of 4MB, then ADSM rounds up to the next higher 4MB partition. For example, if you enter `reduce db 5`, ADSM reduces the recovery log by 8MB.

For example, to reduce the database by 100MB, enter:

```
reduce db 100
```

To reduce the recovery log by 16MB, enter:

```
reduce log 16
```

When you issue a reduce command, the server begins reclaiming space from the last volume to which data was added. If necessary, it relocates data within the defined space. Depending on how much data is relocated, this process can take a long time. When this happens, the reduce operation is run as a background process.

You can query the server for information about the database or recovery log to determine how much free space is available after reduction. For example, after reducing the database by 100MB, the assigned capacity is 96MB, the maximum extension is 100MB, and the maximum reduction is 92MB, as shown in Figure 19.

Available Space (MB)	Assigned Capacity (MB)	Maximum Extension (MB)	Maximum Reduction (MB)	Page Size (bytes)	Total Pages	Used Pages	%Util	Max. %Util
196	96	100	92	4,096	24,576	86	0.3	0.3

Figure 19. Information about the Database

### Step 4: Deleting a Volume from the Database or Recovery Log

After you reduce the database or recovery log, try using the smaller size for a few days. If the maximum utilization percentage does not go over 70%, then you can delete extra volumes from the database or recovery log.

ADSM does not allow you to delete volumes if you do not have sufficient space to store the existing data in the database or recovery log. In addition, ADSM does not allow you to delete the last remaining volume of the database or recovery log.

However, when there is sufficient free space on other volumes, you can delete volumes from the database or recovery log. This process allows you to consolidate data from partially used volumes so that you can use disk space efficiently.

For example, after you reduce the database by 100MB, you determine that you can delete four 24MB volumes from the database.

To delete volumes 1 through 4 from the database, enter the following commands:

```
delete dbvolume vol1
delete dbvolume vol2
delete dbvolume vol3
delete dbvolume vol4
```

To delete the log.1 recovery log volume, enter:

```
delete logvolume log.1
```

When you delete volumes from the database or recovery log, the server moves existing data to available space on other volumes. In this example, as shown in Figure 20, data is moved from volumes 1 through 4 to available space on VOL5. When all data is moved, these volumes are deleted from the server.




Figure 20. Deleting a Database Volume

## Optimizing the Performance of the Database or Recovery Log

When a transaction occurs, the server reads or writes pages to database and recovery log buffer pools. By monitoring the database and recovery log buffer pools, you can optimize the performance of the database and recovery log.

### Database Buffer Pool

When a transaction occurs, one or more database pages may require updating. Each *page* is a 4096-byte block that remains in the buffer pool until space is needed by another page.

An administrator can update the size of the database buffer pool to optimize the I/O performance of the database by updating the BUFPOOLSIZE in the server options

file. “Monitoring the Database Buffer Pool” on page 84 provides information about increasing the size of the database buffer pool.

## Recovery Log Buffer Pool

Because the server supports concurrent client sessions, numerous transactions can occur at the same time. Consequently, the recovery log can be updated continuously. To support concurrent transactions, the recovery log holds transaction log records in a log buffer.

You can update the size of the log buffer to optimize the I/O performance of the recovery log by updating the LOGPOOLSIZ in the server options file. “Monitoring the Recovery Log Buffer Pool” on page 86 provides information about increasing the size of the recovery log buffer pool.

Transaction log records remain in the log buffer until the active buffer becomes full or AD SM forces log records to the recovery log.

When all log records for a transaction are written to the recovery log, the updates are committed to the database. Then the recovery log can be used to roll back the transaction’s updates if there is a system failure or if the transaction does not complete successfully.

## Monitoring the Database Buffer Pool

The *database buffer pool* provides cache storage, which allows database pages to remain in memory for longer periods of time. When database pages remain in cache, the server can make continuous updates to the pages without requiring I/O operations to external storage. While a large database buffer pool can improve server performance, it also requires more memory.

A system administrator can control the size of the database buffer pool by updating the BUFPOOLSIZ option in the server options file. At installation, the database buffer pool is set to 512KB, which equals 128 database pages. For information about setting server options, see *ADSM Installing the Server and Administrative Client*.

### Requesting Information about the Database Buffer Pool

To evaluate whether the current size of the database buffer pool is adequate for database performance, request a detailed database report by entering:

```
query db format=detailed
```

Figure 21 on page 85 displays a detailed database report.






Figure 21. Detailed Information about the Database

Use the following fields to evaluate your current use of the database buffer pool:

**Buffer Pool Pages**

Specifies the number of pages in the database buffer pool. This value is determined by the value set for the BUFPOOLSIZE option in the server options file.

**Total Buffer Requests**

Specifies the cumulative number of requests for database pages since the server was last started or since the last reset of the buffer pool.

**Cache Hit Pct**

Specifies, as a percentage, the number of requests for cached database pages in the database buffer pool that were not read from disk.

A high *cache hit percentage* indicates that the size of your database buffer pool is adequate. If the cache hit percentage drops below 90%, consider increasing the size of the database buffer pool.

**Cache Wait Pct**

Specifies, as a percentage, the number of requests for database pages that had to wait for a buffer to become available in the database buffer pool.

When the cache wait percentage is above 0, increase the size of the database buffer pool.

**Resetting Database Buffer Pool Statistics**

To gather statistics on database use, reset the buffer pool statistics on a regular basis and chart the results.

Initially, you might want to monitor the database twice a day. Later, when most client nodes have been registered to the server, you can reset statistics on a weekly basis.

For example, to reset the database buffer pool, enter:

```
reset bufpool
```

## Monitoring the Recovery Log Buffer Pool

The *recovery log buffer pool* is used to hold new transaction records until they can be written to the recovery log. The size of the recovery log buffer pool can affect how frequently the server forces records to the recovery log.

A system administrator can control the size of the recovery log buffer pool by updating the LOGPOOLSIZE option in the server options file. At installation, the default setting is 128KB, which equals 32 recovery log pages. For information about setting server options, see *ADSM Installing the Server and Administrative Client*.

### Requesting Information about the Recovery Log Buffer Pool

To determine how the buffer pool size affects recovery log performance, query the recovery log for a detailed report by entering:

```
query log format=detailed
```

Figure 22 displays a detailed report.

```
Available Space (MB): 12
Assigned Capacity (MB): 12
Maximum Extension (MB): 0
Maximum Reduction (MB): 8
  Page Size (bytes): 4,096
    Total Pages: 3,072
      Used Pages: 227
        %Util: 7.4
          Max. %Util: 69.6
            Physical Volumes: 1
              Log Pool Pages: 32
                Log Pool Pct. Util: 6.25
                  Log Pool Pct. Wait: 0.00
```

Figure 22. Detailed Information about the Recovery Log

Use the following fields to optimize the log buffer pool size for your installation:

#### Log Pool Pages

Specifies the number of pages in the recovery log buffer pool. This value is determined by the value set for the LOGPOOLSIZE option in the server options file.

#### Log Pool Pct. Util

Specifies, as a percentage of all recovery log buffer pool pages, the number of pages used to write changes to the recovery log after a transaction is committed.

A low *log pool percent utilization* indicates that the size of your recovery log buffer pool is adequate. As this number grows larger, consider increasing the size of the recovery log buffer pool.

**Log Pool Pct. Wait**

Specifies, as a percentage of all recovery log buffer pool pages, the number of requests for a page that is not available because all pages are waiting to write to the recovery log.

If the *log pool percentage wait* value is greater than zero, increase the size of the recovery log buffer pool.



---

## Chapter 5. Managing Licensing, Privilege Classes, and Registration

This section provides the information necessary for a system administrator to control authorization and access to the server. The sections listed in the following table begin at the indicated pages.

Section	Page
<b>Tasks:</b>	
Managing ADSM licenses	89
Ensuring client/server authentication	92
Registering administrators or updating information	93
Granting administrative authority	94
Revoking or reducing administrative authority	98
Managing administrator access	100
Managing client node registration	102
Registering an application programming interface to the server	110
Managing client node access	105
Requesting information about client nodes or file spaces	107
Deleting client data and client nodes from the server	109

Most tasks presented in this chapter can be performed using either the graphical user interface or the command-line interface. Table 4 on page 33 shows whether a task can be performed on the graphical user interface, the command line interface, or both.

For information about issuing ADSM commands, refer to *ADSM Administrator's Reference*. For help performing a task on the graphical user interface, see the procedure described in "Using Online Help" on page 15.

---

### Managing ADSM Licenses

Task	Required Privilege Class
<ul style="list-style-type: none"><li>• Register licenses</li><li>• Audit licenses</li><li>• Schedule automatic license audits</li></ul>	System
Display license information	Any administrator

If an ADSM system exceeds the terms of its license agreement, one of the following occurs:

- The server issues a warning message indicating that it is not in compliance with the licensing terms.
- Operations fail because the server is not licensed for specific features.

For details about either case, see “License Compliance” on page 92. In either case, you must contact your IBM account representative or authorized reseller to modify your agreement.

## Licensed Features

The base ADSM MVS server license supports an unlimited number of administrative clients, one OS/2 backup-archive client, and a specified set of removable media devices.

**Note:** In this licensing section, the term *client* is used to refer to backup-archive clients, unless otherwise noted.

You must register a new license if you want to make any of the following changes to your license agreement:

- Add support for additional clients. The base license allows for one OS/2 backup-archive client. If you want to add clients in an environment other than OS/2, you must register a new license for that feature also (see the next item in this list). See “Additional Clients.”
- Add support for clients in environments other than OS/2. The base license allows only for backup-archive clients on OS/2. See “Clients Other Than OS/2.”
- Add support for storage devices not covered by the existing agreement. See “Device Support Modules 3-4” on page 91.
- Add a secondary server. See “Secondary Server Attachment” on page 91.

A REGISTER LICENSE command adds a license statement to the end of the server options file. For this to happen, the options file must be specified in the OPT DD statement of the server startup JCL with a DISP=MOD parameter. You can also register licenses by specifying them directly in the options file. See *ADSM Installing the Server and Administrative Client* for details.

### Additional Clients

You register the server to support a specified number of clients beyond the one OS/2 backup-archive client supported by the base license. Those additional clients can be in any environment for which your system is licensed (see “Clients Other Than OS/2”).

For example, to register three additional clients, enter:

```
register license clients 3
```

If you register more clients than your server is licensed to support, the server issues a warning message. However, operations continue normally.

### Clients Other Than OS/2

You can obtain licenses for environment support features that allow the server to support clients other than OS/2

Environment support features are:

#### **DESKTOP**

DOS, Macintosh, NetWare, or Windows

## UNIX

AIX, AT&T, DEC ULTRIX, HP-UX, SCO UNIX 386, SCO Open Desktop, SINIX-Z, SunOS or Solaris, or UNIX

## OPENEDITION

Open Edition for MVS

## SPACEMGMT

HSM clients on AIX, HP-UX, SunOS or Solaris

To register more than one environment support feature, issue a separate REGISTER LICENSE command for each feature. For example, to allow AIX clients and HSM clients, issue:

```
register license unix
register license spacemgmt
```

## Device Support Modules 3-4

You can obtain licenses for device support modules that allow the server to support a variety of storage devices. Device support modules for storage devices are numbered 1 through 4, and each module includes all devices supported by any lower-numbered module. For example, Device Support Module 4 supports any device supported by Device Support Modules 1, 2, and 3. The MVS base license includes Device Support Modules 1 and 2.

To let the server attach storage devices in Device Support Module 3, enter:

```
register license devicemod3
```

Any attempt to define a library or drive that requires a device support module fails if the module is not registered. If you try to mount a volume requiring a library or drive that is not licensed, the operation also fails.

The *ADSM Licensed Program Specifications and License Information* list the devices and libraries supported by each device support module. However, device support is continually expanded. For current information about supported devices, check with IBM or your authorized reseller, or call the IBM Information Support Center at 1-800-IBM-3333 and ask for STAR 20.

## Secondary Server Attachment

You can obtain a license for attaching a secondary server to a library. For example, if you have a license for Device Support Module 4, you can get a license that lets you attach a secondary server to a library in that module. Register that license by entering:

```
register license secondaryserverattach
```

## License Compliance

If license terms change (for example, a new license is specified for the server), the server conducts an audit to determine if the current server configuration conforms to the license terms.

The server also periodically audits compliance with the license terms. The results of this audit are used to check and enforce license terms. If 30 days have elapsed since the previous license audit, the administrator cannot cancel the audit.

The number of client nodes for which a server is licensed is enforced when the server is in open registration mode. If the terms of the license are violated by the addition of another registered node, the server issues a warning message stating that it is out of compliance.

If the server is not licensed to support a type of client (environment support) or device (device support module), server operations fail when you try to use the client or device. If one or more of the features or device support modules are licensed on the server, you receive error messages if you exceed your license terms.

## Licensing Example

You must issue a separate REGISTER LICENSE command for each feature or device support module. For example, to license a server for one or more features or device support modules:

1. Obtain the additional licenses for the features or device support modules from your IBM account representative or authorized reseller.
2. Issue a REGISTER LICENSE command for each feature or device support module.

## Monitoring Licenses

An administrator can monitor license compliance by:

### Auditing licenses

Use the AUDIT LICENSES commands or the GUI to compare the current configuration with the current licenses.

### Displaying license information

Use the QUERY LICENSE command or the GUI to display details of your current licenses and determine licensing compliance.

### Scheduling automatic license audits

Use the SET LICENSEAUDITPERIOD command or the GUI to specify the number of days between automatic audits.

---

## Ensuring Client/Server Authentication

Task	Required Privilege Class
Set password authentication Set password expiration	System

ADSM provides client/server authentication to validate that administrative clients or client nodes are communicating with an authorized server, and that the server is communicating with registered administrators and client nodes. ADSM ensures



client/server authentication by requiring a password from each administrative client or client node registered with a server.

You can ensure client/server authentication by:

- Requiring users to enter a password to access the server
- Requiring users to change their passwords regularly

## Setting Client Password Authentication

At installation ADSM automatically sets password authentication on. With password authentication set to on, all users must enter a password when accessing the server. ADSM maintains a list of registered client passwords for authentication.

To allow administrators and client nodes to access ADSM without entering a password, you must set password authentication to off by issuing the SET AUTHENTICATION command.

**Warning:** Setting password authentication off, reduces data security.

## Setting User Password Expiration

You can set a password expiration period for administrators and client node users by issuing the SET PASSEXP command. The valid password expiration period is from 1 to 9999 days. At installation, ADSM sets a password expiration of 90 days.

When an administrator or client node user is first registered to the server, ADSM begins tracking the password expiration period. For example, when using the installed value, users must change their password within 90 days of the password expiration period. For example, when using the installed value, users must change their password within 90 days after they are registered to the server.

If a user password is not changed within this period, the server prompts the user to change the password the next time the user attempts to access the server.

---

## Registering Administrators or Updating Information

Task	Required Privilege Class
Register an administrator or update information about other administrators	System
Update information about yourself	Any administrator

To register an administrator, specify a user ID and password for the user. You also can provide contact information such as the user name and telephone number. This contact information is displayed when administrators query for administrator information (Format=Detailed).

To register the administrator with a user ID of DAVEHIL and the password of *birds*, enter the REGISTER ADMIN command:

```
register admin davehil birds contact='backup team'
```

After users are registered as administrators, they can display ADSM server information from any computer on which the administrative client program is installed.

**Note:** At installation, the server console is defined with a special user ID, which is named SERVER\_CONSOLE. This name is reserved and cannot be used by another administrator. At installation, the SERVER\_CONSOLE user ID is granted system privilege so that other administrators can be registered and granted system privilege.

Another administrator with system privilege can revoke or grant new privileges to the SERVER\_CONSOLE user ID. However, you cannot update, lock, rename, or remove the SERVER\_CONSOLE user ID from ADSM.

See the *ADSM Installing the Server and Administrative Client* for information on using the SERVER\_CONSOLE user ID.

If administrator DAVEHIL forgets his password, a system administrator can update the password so that DAVEHIL can access the ADSM server again. To change the password for administrator ID DAVEHIL to *ganymede*, enter:

```
update admin davehil ganymede
```

**Note:** The SERVER\_CONSOLE administrator's user ID and contact information cannot be updated.


---

## Granting Administrative Authority

Task	Required Privilege Class
Grant authority to other administrators	System

After administrators have been registered, they can make queries and request command-line help. To perform other ADSM functions, they must be granted authority by being assigned one or more administrative privilege classes.

This section describes the privilege classes, which are illustrated in the following figure. An administrator with system privilege can perform any ADSM function. Administrators with policy, storage, operator, or analyst privileges can perform a subset of ADSM functions.



Whenever you issue a subsequent GRANT AUTHORITY command, ADSM adds authority to the administrator; it does not override existing authority.

## System Privilege

An administrator with *system privilege* can perform any ADSM administrative task. To perform the following tasks you must have system privilege:

- Register or remove administrators
- Grant or revoke all levels of administrative authority
- Lock or unlock administrators from the server
- Rename administrators or update administrator information
- Define or delete policy domains and storage pools
- Import or export data from the server
- Cancel administrative background processes
- Set operating parameters for the server
- Perform license audits

In addition, an administrator with system privilege can grant other administrators system privilege by specifying an administrator user ID and the system privilege class.

For example, to grant the system privilege class to administrator KACZ, enter:

```
grant authority kacz classes=system
```

## Unrestricted Policy Privilege

An administrator with *unrestricted policy privilege* has the authority to manage the backup and archive services for client nodes assigned to any policy domain. When new policy domains are defined to the server, an administrator with unrestricted policy privilege is automatically authorized to manage the new policy domains.

An administrator with unrestricted policy privilege can:

- Register client nodes in any policy domain

- Manage any client node access to the server
- Delete any client node files from storage pools
- Manage policy objects within any policy domain
- Manage schedules, that automatically back up or archive files
- Associate client nodes to schedules defined in the same policy domain

To grant unrestricted policy privilege to administrator SMITH, enter:

```
grant authority smith classes=policy
```

However, SMITH cannot copy, define, or delete policy domains.

## Restricted Policy Privilege

An administrator with *restricted policy privilege* is authorized to issue commands only for the policy domains to which they have been authorized.

An administrator with restricted policy privilege can:

- Register a client node to an authorized policy domain
- Manage access for client nodes assigned to an authorized policy domain
- Delete files from storage pools for client nodes in authorized policy domains
- Manage policy objects in authorized policy domains
- Manage backup or archive schedules in authorized policy domains
- Associate schedules to client nodes assigned to an authorized policy domain

To grant restricted policy privilege over the policy domain named ENGPOOLDOM, to administrator SMITH enter:

```
grant authority smith domains=engpoldom
```

However, DAVEHIL (as established in “Registering Administrators or Updating Information” on page 93) cannot copy a management class from the engineering policy domain to the standard policy domain because the administrator does not have authority over the policy domain named STANDARD.

## Unrestricted Storage Privilege

An administrator with *unrestricted storage privilege* has the authority to manage the database, recovery log, and all storage pools.

An administrator with unrestricted storage privilege can:

- Define volumes to the database or recovery log
- Extend or reduce the size of the database or recovery log
- Create mirrored copy sets of the database or recovery log
- Delete volumes from the database or recovery log
- Manage disk and tape device classes
- Define volumes to any disk or tape storage pools
- Move data from a storage pool to any other storage pool
- Delete volumes from any storage pool
- Audit volumes belonging to any storage pool

To grant unrestricted storage privilege to administrator COYOTE, enter:

```
grant authority coyote classes=storage
```

## Restricted Storage Privilege

Administrators with *restricted storage privilege* can issue a subset of the storage commands only for the storage pools for which they have been authorized. They do not have authority to manage the database or recovery log.

An administrator with restricted storage privilege can:

- Define volumes to authorized disk or tape storage pools
- Move data from a volume to another volume in an authorized storage pool
- Delete volumes from an authorized storage pool
- Audit volumes belonging to an authorized storage pool

To grant restricted storage privilege for existing storage pools beginning with the name ADSM to administrator HOLLAND, enter:

```
grant authority holland stgpools=adsm*
```

Administrator HOLLAND is restricted to managing storage pools beginning with ADSM that existed at when the command was first issued; for example, ADSM.BFS.TAPE1, ADSM.BFS.TAPE2, and so on. HOLLAND is not authorized to manage any *new* storage pools that are defined after authority has been granted.

To add a new storage pool, ADSM.BFS.TAPEX, to HOLLAND's authority, enter:

```
grant authority holland stgpools=adsm.bfs.tapex
```

## Operator Privilege

Administrators with *operator privilege* control the immediate operation of the ADSM server and the availability of storage media.

An administrator with operator privilege can:

- Disable the server to prevent clients from accessing the server
- Enable the server for access by clients
- Cancel client/server sessions
- Vary disk volumes on or off line to perform maintenance
- Reset the error status for tape volumes
- Manage tape mounts
- Halt the server, when necessary

To grant operator privilege to administrator HOLLAND, enter:

```
grant authority holland classes=operator
```

Now, HOLLAND has operator privilege so that he can manage tape operations, as well as having restricted storage privilege over tape storage pools.

## Analyst Privilege

An administrator with *analyst privilege* can issue commands that reset the counters that track server statistics.

To grant analyst privilege to administrator MARYSMITH, enter:

```
grant authority marysmith classes=analyst
```

---

## Revoking or Reducing Administrative Authority

Task	Required Privilege Class
Revoke or reduce administrative privilege classes	System

You can remove all or part of another administrator's authority. You can reduce administrative authority by revoking:

- The system privilege class from an administrator and then granting one or more of the following privilege classes: storage, policy, operator, or analyst
- Unrestricted policy or storage privileges and then granting restricted policy or storage privileges to the administrator
- Authority over some policy domains from an administrator with restricted policy privilege
- Authority over some storage pools from an administrator with restricted storage privilege

## Revoking All Administrative Privilege Classes

You can revoke all administrative privilege classes from another administrator so that the latter administrator is only allowed to perform those functions that can be performed by any administrators.

To revoke all administrative privilege classes from an administrator, identify the administrator user ID, but do not specify any privilege classes, policy domains, or storage pools. For example, to revoke both the storage and operator privilege classes from administrator HOLLAND issue the REVOKE AUTHORITY command:

```
revoke authority holland
```

HOLLAND can still query the server for information, but he can no longer issue administrative commands.

## Revoking One or More Administrative Privilege Classes

You can revoke part of an administrator's authority by specifying the administrator's user ID and one or more privilege classes.

Assume that administrator MARYSMITH no longer has the time to manage the database, recovery log, or disk and tape storage pools. To remove part of her administrative authority by revoking the storage privilege class, enter:

```
revoke authority marysmith classes=storage
```

## Reducing System Authority

You can reduce the authority of another system administrator by revoking the system privilege class and granting one or more of the remaining privilege classes.

For example, to reduce the authority of administrator SERVER\_CONSOLE to the operator privilege class do the following:

1. Revoke the system privilege class from administrator SERVER\_CONSOLE by entering:

```
revoke authority server_console classes=system
```

2. Grant operator privilege class to administrator SERVER\_CONSOLE by entering:

```
grant authority server_console classes=operator
```

## Reducing Unrestricted Policy to Restricted Policy Privilege

You can reduce an administrator's unrestricted policy privilege by revoking the policy privilege class and then granting restricted policy privilege over specified policy domains.

For example, to reduce administrator DSMITH's unrestricted policy privilege to restricted policy privilege for the STANDARD policy domain, do the following:

1. Revoke the unrestricted policy privilege class from DSMITH by entering:

```
revoke authority dsmith classes=policy
```

2. Grant restricted policy privilege over the STANDARD policy domain by entering:

```
grant authority dsmith domains=standard
```

## Reducing Unrestricted Storage to Restricted Storage Privilege

You can reduce an administrator's unrestricted storage privilege by revoking the storage privilege class and granting restricted storage privilege over specified storage pools.

For example, to reduce administrator COYOTE from unrestricted storage privilege to restricted storage privilege for the storage pools named BACKUPPOOL and ARCHIVEPOOL, do the following:

1. Revoke the unrestricted storage privilege class from COYOTE by entering:

```
revoke authority coyote classes=storage
```

2. Grant restricted storage privilege over the BACKUPPOOL and ARCHIVEPOOL storage pools by entering:

```
grant authority coyote stgpools=backuppools,archivepool
```

## Reducing Restricted Policy or Storage Privilege

You can reduce restricted policy privilege by revoking authority over previously authorized policy domains. You can also reduce restricted storage privilege by revoking authority over previously authorized storage pools.

For example to reduce administrator COYOTE's, restricted storage privilege, revoke authority for the storage pool named ARCHIVEPOOL by entering:

```
revoke authority coyote stgpools=archivepool
```

---

## Managing Administrator Access

A system administrator can control access to the server by renaming an administrator, removing an administrator, or by locking and unlocking an administrator from the server. This section describes these methods and gives examples.

Task	Required Privilege Class
Rename an administrator user ID Remove other administrators from the server Temporarily prevent other administrators from accessing the system	System privilege
Display administrator information	Any administrator



## Renaming an Administrator

You can rename an administrator ID when an employee wants to be identified by a new ID, or you want to assign an existing administrator ID to another person. You cannot rename an administrator ID to one that already exists on the system.

For example, if administrator HOLLAND leaves your organization, you can assign administrative privilege classes to another user by completing the following steps:

1. Assign HOLLAND's user ID to WAYNESMITH by issuing the RENAME ADMIN command:

```
rename admin holland waynesmith
```

By renaming the administrator ID, you remove HOLLAND as a registered administrator from the server. In addition, you register WAYNESMITH as an administrator with the password, contact information, and administrative privilege classes previously assigned to HOLLAND.

2. Change the password to prevent the previous administrator from accessing the server by entering:

```
update admin waynesmith new_password contact="development"
```

**Note:** The administrator SERVER\_CONSOLE cannot be renamed.

## Removing Administrators

You can remove other administrators from the server so that they no longer have access to administrator functions. For example, to remove registered administrator ID SMITH, enter:

```
remove admin smith
```

### Notes:

1. ADSM does not allow you to remove the last system administrator from the system.
2. The administrator user ID named SERVER\_CONSOLE cannot be removed.

## Locking and Unlocking Administrators from the Server

To temporarily prevent other administrators from accessing the system, a system administrator can lock out other administrators from ADSM.

For example, you can lock out an administrator when an employee takes a leave of absence from your business. You can lock out the administrator user ID MARYSMITH from the server user ID by issuing the LOCK ADMIN command:

```
lock admin marysmith
```

When the employee returns, any system administrator can unlock the administrator ID. For example, you can unlock MARY SMITH's user ID by issuing the UNLOCK ADMIN command:

```
unlock admin marysmith
```

MARYSMITH can now access ADSM to complete administrative tasks.

**Note:** You cannot lock or unlock the SERVER\_CONSOLE administrator user ID from the server.

### Requesting Information about Administrators

Any administrator can query the server to view administrator information. You can also query all administrators authorized with a specific privilege class.

For example, to query the system for a detailed report on administrator ID DAVEHIL, issue the QUERY ADMIN command:

```
query admin davehil format=detailed
```

Figure 23 displays a detailed report.

```
Administrator Name: DAVEHIL
Last Access Date/Time: 02/09/1994 19:49:46
Days Since Last Access: 1
Password Set Date/Time: 02/08/1994 19:49:31
Days Since Password Set: 1
    Locked?: No
    Contact: backup team
System Privilege:
Policy Privilege: ENGPOLDOM
Storage Privilege:
Analyst Privilege:
Operator Privilege:
Registration Date: 02/09/1994 19:00:00
Registering Administrator: REES
```

Figure 23. A Detailed Administrator Report

## Managing Client Nodes

Task	Required Privilege Class
Set registration to open or closed	System
Register client nodes to any policy domain	System or unrestricted policy
Register client nodes to specific policy domains	System, unrestricted policy, or restricted policy for those domains
Update, rename, lock, or unlock any client nodes	System, unrestricted policy

Task	Required Privilege Class
Update, rename, lock, or unlock client nodes assigned to specific policy domains	System, unrestricted policy, or restricted policy for those domains
Request information about client nodes or file spaces	Any administrator
Delete any file space from storage pools	System or unrestricted policy
Delete file spaces defined for client nodes assigned to specific policy domains	System, unrestricted policy, or restricted policy for those domains
Remove any client nodes	System or unrestricted policy
Remove client nodes assigned to specific policy domains	System, unrestricted policy, or restricted policy for those domains

Managing client node registration includes:

- Setting client node registration to open or closed
- Registering client nodes
- Updating client node information
- Managing client node access
- Requesting information about client nodes and file spaces
- Deleting file spaces and client nodes

## Setting Client Node Registration

Before a user can begin requesting backup or archive services from an ADSM server, the workstation must be registered with a server.

ADSM provides two methods for registering workstations with an ADSM server:

### Open registration

Users register their own workstations as client nodes with the server.

### Closed registration

An administrator registers each workstation as a client node with the server.

At installation, registration is set to closed.

**Note:** You can use the SET REGISTRATION command to change registration at any time. Existing registered client nodes are not affected by changes to the registration process.

### Open Registration

With open registration, users can register their own workstations with the server. When a user attempts to access the server from an unregistered workstation, the server prompts the user for a password and contact information and registers the workstation as a client node with the server. On UNIX systems, only the root user can register a workstation as a client node with the server.

ADSM sets the following defaults:

- Assigns each client node to the policy domain STANDARD.
- Allows each client user to choose whether or not to compress files. A user with root authority can define whether compression is used by entering the appropriate value for the COMPRESSION option in the **dsm.opt** client options file.

- Allows each client node user to delete archived copies (but not backed up files) from storage pools.

To change any of these defaults after the client node has been registered, you can update client node registration information as described in “Updating Client Node Information” on page 105.

## Closed Registration

With closed registration, a system administrator, an administrator with unrestricted policy privilege, or an administrator with restricted policy privilege can define the following information for each client node:

- The workstation node name. UNIX users must provide the value returned by the HOSTNAME command to the administrator.
- The user password.
- The policy domain to which the client node is assigned. The policy domain contains the policy objects that control how ADSM manages user backup and archive data.
- Whether the user can compress files at the workstation before sending them to the server. You can select one of three options to specify whether the client program should:
  - Compress files before sending them to the server. Time is saved when objects are compressed before they are backed up or archived to the server. In addition, server storage space is saved.
  - Send files that are not compressed to the server  
 Compression can affect ADSM throughput and require more workstation memory. Typically, a workstation with a slow processor connected to the server on a high-speed transmission line does not benefit from compression. To optimize performance or to ease memory constraints at the workstation, an ADSM administrator can restrict file compression.
  - Use the value defined for the COMPRESSION option to determine whether or not to compress objects being sent to the server  
 The COMPRESSION option can be set in the application configuration file or in the client system options file.
- Whether the user is allowed to delete backed up or archived files from storage pools, by using the DSMC DELETE FILESPACE or DSMC DELETE ARCHIVE command.

If users are not allowed to delete archived or backed up files, an administrator with system or policy privilege must delete any files associated with the client nodes from storage pools. See “Deleting File Spaces and Client Nodes” on page 109 for information on deleting files from storage pools.

## Registering Client Nodes

Authorized administrators can register client nodes with the server. If registration is open, then users can register their own workstation with the server. See “Setting Client Node Registration” on page 103 for information about setting open or closed registration.

To register a client node, you define:

- The workstation node name
- The user password and contact information
- Whether the user client program should compress files
- Whether the user is allowed to delete backed up files from storage
- Whether the user is allowed to delete archived files from storage

For example, you want to register three workstations from your engineering department to the server. Because your engineers have unique storage management needs, you want to assign them to the *ENGPOLDOM* policy domain. Before you can assign client nodes to a policy domain, you must define the policy domain. See Chapter 6, “Managing Policies” on page 113 for information on defining new policy domains.

For each workstation, define a node name, password, and contact information. Set file compression on so that, when users back up or archive files, the client node attempts to compress the files before transmitting them to storage pools. Finally, allow these engineers to delete backed up or archived files from storage pools.

If you are administering the server from an administrative client, you can use the macro facility to register more than one client node at a time.

For this example, create a macro file named *REGENG.MAC*. Enter the *REGISTER NODE* commands in the macro file by using your preferred workstation editor:

```
register node ssteiner choir contact='department 21' -
domain=engpoldom compression=yes archdelete=yes cackdelete=yes
register node carolh skiing contact='department 21, second shift' -
domain=engpoldom compression=yes archdelete=yes cackdelete=yes
register node mab guitar contact='department 21, third shift' -
domain=engpoldom compression=yes archdelete=yes backdelete=yes
```

After you create the macro file, invoke the macro from an administrative client by issuing the *MACRO* command:

```
macro regeng.mac
```

ADSM processes the commands in the macro file and registers *ssteiner*, *carolh* and *mab* client nodes with the server.

## Managing Client Node Access

You can control client node access to ADSM by updating or renaming client nodes or by locking and unlocking a client nodes from the server. This section describes these methods and gives examples.

### Updating Client Node Information

You can update the following client node registration information:

- The user password or contact information
- The policy domain to which the client node is assigned

**Note:** An administrator with restricted policy privilege must be authorized to the current policy domain and to the new policy domain.

- Whether file compression is required
- Whether users can delete backed up or archived files from storage pools

For example, if a user registers the workstation during open registration as administrator TOMC, ADSM:

- Assigns TOMC to the policy domain named STANDARD
- Allows TOMC to define whether or not to compress files from storage pools
- Allows TOMC to delete archived, but not backed up, files from storage pools

Novice ADSM users can be restricted from deleting archived files from storage pools. You can update the node registration for the TOMC client node by issuing the UPDATE NODE command:

```
update node tomc archdelete=no
```

TOMC is now restricted from deleting archived files from storage pools. Only an administrator with system or policy privilege can delete TOMC's archived files.

### Renaming Client Nodes

You can rename a client node if the workstation network name or host name is changed. For example, with UNIX systems, users define their ADSM node named based on the value returned by the HOSTNAME option. When users access the server, their ADSM user IDs match the host name of their workstations.

If the host name changes, you can update a client node user ID to match the new host name. For example, you can rename CAROLH to ENGNODE by issuing the RENAME NODE command:

```
rename node carolh engnode
```

The client node named ENGNODE retains contact information and access to backup and archive data. In addition, all files backed up or archived by CAROLH now belong to the client node named ENGNODE.

### Locking and Unlocking Client Nodes

You can lock a client node from accessing the server, and you can unlock the client node to allow a user to reaccess the server. For example, you can lock the client node named MAB from the server by issuing the LOCK NODE command:

```
lock node mab
```

When the client node is locked from the server, a user is prevented from accessing the server to perform ADSM functions, such as backing up or archiving files to the server, or restoring or retrieving files from the server.

You can allow client node MAB to reaccess the server by issuing:

```
unlock node mab
```

## Requesting Information about Client Nodes or File Spaces

After client nodes are registered with the ADSM server, users have the option of defining file spaces on their workstation. A *file space* name identifies a group of files that are stored as a logical unit in data storage.

On client systems such as OS/2 or DOS, a file space name identifies a logical partition, such as the volume label of a disk drive. For example, a volume with the label XYZ is a different file space than a volume with the label ABC.

On client systems such as AIX or SunOS, a file space name identifies a file system or file space defined by a user with the VIRTUALMOUNTPOINT option. With the VIRTUALMOUNTPOINT option, users can define a virtual mount point for a file system to back up or archive files beginning with a specific directory or subdirectory. For information on the VIRTUALMOUNTPOINT option, refer to the appropriate *ADSM Using the Backup-Archive Client*.

### Requesting Information about Client Nodes

You can request information about client nodes. For example, as a policy administrator, you might query the server about all client nodes assigned to the policy domains for which you have authority. Or you might query the server for detailed information about one client node.

**Client Nodes Assigned to Specified Policy Domains:** You can display information about client nodes assigned to specific policy domains. For example, to query the server to view information about any client nodes assigned to the policy domains named STANDARD and ENGPOLDOM, issue the QUERY NODE command:

```
query node * domain=standard,engpoldom
```

Figure 24 displays a standard report.

Node Name	Platform	Policy Domain Name	Days Since Last Access	Days Since Password Set	Locked?
DEBBYG	DOS	STANDARD	2	12	No
ENGNODE	AIX	ENGPOLDOM	<1	1	No
HTANG	OS/2	STANDARD	4	11	No
MAB	AIX	ENGPOLDOM	<1	1	No
PEASE	AIX	STANDARD	3	12	No
SSTEINER	(?)	ENGPOLDOM	<1	1	No

Figure 24. Requesting Client Node Information

**A Specific Client Node:** You can display detailed information about specific client nodes. This is useful when you want to review the registration parameters defined for users. For example, you can display a detailed report for the client node named PEASE, by issuing the QUERY NODE command:

```
query node pease format=detailed
```

The following figure displays the contents of a detailed report for the client named PEASE.

```
Node Name: PEASE
Platform: AIX
Policy Domain Name: STANDARD
Last Access Date/Time: 02/21/1994 10:58:36
Days Since Last Access: 3
Password Set Date/Time: 02/09/1994 10:02:00
Days Since Password Set: 12
  Locked?: No
  Contact:
Compression: Yes
Archive Delete Allowed?: No
Backup Delete Allowed?: No
Registration Date: 02/09/1994 10:02:00
Registering Administrator: REES
Last Communication Method Used: Tcp/Ip
Bytes Received Last Session: 1,719
Bytes Sent Last Session: 602
Duration of Last Session (sec): 184.63
Pct. Idle Wait Last Session: 99.69
Pct. Comm. Wait Last Session: 0.00
Pct. Media Wait Last Session: 0.00
```

## Requesting File Space Information

You can display file space information in order to:

- Identify file spaces defined to each client node, so that you can delete each file space from the server before removing the client node from the server
- Monitor the actual space consumed on workstation's disks
- Monitor whether backups are completing successfully for the file space
- Determine the date and time of the last backup

You display file space information by identifying the client node name and file space name. Be aware that file space names are case-sensitive. You can request a standard or detailed report.

For example, you can query the server for information about file spaces defined for the client node named PEASE, by issuing the QUERY FILESPACE command:

```
query filespace pease *
```



The figure below displays the output from this command. It shows that node ID PEASE has defined three file spaces on his AIX workstation, which has been registered as a client node named PEASE. It also shows that node ID PEASE is running the *JFS* file system on his AIX workstation. Finally, this report shows how much space is available in each file space on the node ID PEASE's workstation and how much space is being used.

Node Name	Filespace Name	Platform	Filespace Type	Capacity (MB)	%Util
PEASE	/home/pease/dir	AIX	JFS	196.0	91.7
PEASE	/home/pease/dir1	AIX	JFS	328.0	81.0
PEASE	/home/pease/dir2	AIX	JFS	46.9	96.0

Figure 25. Report Received from the Query Filespace Command

## Deleting File Spaces and Client Nodes

You can delete a client node from a server. First, however, you must delete any client backup, archive, and client migrated data from storage pools by deleting the filesystem belonging to the node.

### Deleting a File Space

You may want to delete a file space when:

- Users are not authorized to delete backed up or archived files in storage pools

The authority to delete backed up or archived files from server storage is set when a client node is registered to the server. See “Setting Client Node Registration” on page 103 and “Registering Client Nodes” on page 104 for information on allowing users to delete files in storage pools.

For example, node ID PEASE does not have the authority to delete files that he no longer needs in the file space named */home/pease/dir2*. An administrator must delete the files for him by issuing the `DELETE FILESPACE` command:

```
delete filesystem pease /home/pease/dir2 type=archive
```

- You want to remove a client node from the server

You must delete a user's files from storage pools before you can remove a client node. For example, to delete all files stored in any file spaces defined to the client node ID DEBBYG, enter:

```
delete filesystem debbyg * type=any
```

- You want to delete files belonging to a specific owner

For client nodes that support multiple users, such as UNIX, a file owner name is associated with each file on the server. The owner name is the user ID of the operating system, such as the UNIX user ID. When you delete a file space belonging to a specific owner, only files that have the specified owner name in the file space are deleted.

## Removing Client Nodes

Before you can remove a client node from the server, all backed up or archived files belonging to the client node must be deleted from storage pools. See “Deleting a File Space” on page 109 for information on deleting backed up or archived files from storage pools.

For example, you can remove the client node ID DEBBYG from the server, by issuing the REMOVE NODE command:

```
remove node debbyg
```

A client node that has been removed from the server can no longer back up, archive, or migrate files to ADSM.

---

## Registering an Application Programming Interface to the Server

Workstation users can begin requesting services from an ADSM server by using an application that uses the ADSM application programming interface (API). An administrator must register the workstation as a client with an ADSM server. After the workstation is registered with a server, it can begin to back up, archive, restore, and retrieve objects by using the application’s interface.

To register an API to the server obtain the following information from API users:

- Node name of client
- Initial password, if a password is required
- Contact information such as the user name, user ID, and telephone number

To register an application by using the ADSM API as a client node, define the following:

- Workstation node name
- User password and client information
- Compression status of the client
- Policy domain to which your workstation belongs

## Understanding How the Compression Option is Set

For applications that use the ADSM API, compression can be determined by:

- An administrator during registration who can:
  - Require that files are compressed by the client before they are sent to an ADSM server
  - Restrict files from being compressed by the client
  - Allow the application or client user to determine the compression status

- The client options file. If an administrator does not set compression on or off, ADSM checks the compression status set in the client options file. The client options file is required, but the API user configuration file is optional.
- One of the object attributes. When an application sends an object to the server, some object attributes can be specified. One of the object attributes is a flag that indicates whether or not the data has already been compressed. If the application turns this flag on during either a backup or an archive operation, then ADSM does not compress the data a second time. This process overrides what the administrator sets during registration.

## Understanding How the File Deletion Option is Set

For applications using the ADSM API, the file deletion option can be set by:

- An administrator during registration

If an administrator does not allow the file deletion, then an ADSM administrator must delete any objects or file spaces associated with the workstation from data storage.

If an administrator allows file deletion, then ADSM checks the client options file.

- An application using the ADSM API deletion program calls

If the application uses the **dsmDeleteObj** or **dsmDeleteFS** program call, then objects or files are marked for deletion when the application is executed.



---

## Chapter 6. Managing Policies

ADSM policies control how and when user files are backed up and archived to server storage and how user files are migrated to server storage.

The sections listed in the following table begin at the indicated pages.

Section	Page
<b>Concepts:</b>	
Policy operations	113
Policy objects	115
Management classes	116
Expiration processing	119
File eligibility for policy operations	120
How client migration works with backup and archive	122
<b>Tasks:</b>	
Using the standard storage management policies	122
Creating your own storage management policies	123
Defining a policy domain	126
Defining a policy set	127
Defining a management class	127
Defining a backup copy group	128
Defining an archive copy group	132
Assigning a default management class	134
Validating and activating policy sets	134
Starting expiration processing	135
Querying policy objects	136
Deleting policy objects	138

Most tasks presented in this chapter can be performed by using either the graphical user interface or the command-line interface. Table 5 on page 36 shows whether a task can be performed on the graphical user interface, the command-line interface, or both.

For information about issuing ADSM commands, refer to *ADSM Administrator's Reference*. For help performing a task on the graphical user interface, see the procedure described in "Using Online Help" on page 15.

---

### Policy Operations

ADSM policies govern the following operations, which are discussed in this section:

- Backup and restore
- Archive and retrieve
- Client migration and recall

## Backup and Restore

To guard against the loss of information, ADSM can copy files, subdirectories, and directories to media controlled by ADSM. Backups can be controlled by administrator-defined policies, or users can request backups of their own data. ADSM provides two types of backup:

### Incremental backup

The backup of files that have changed since their last backup. ADSM does the following incremental backup processing:

1. Compares a file on the client workstation with its copy in storage.
2. Determines if the file has changed since the last backup.
3. Backs up files that have been changed and that meet policy eligibility requirements.

### Selective backup

Backs up files that the user specifies and that meet any requirements defined by the applicable backup copy group.

When a user restores a backup version of a file, ADSM sends a copy of the file to the client node. The backup version remains in ADSM storage.

If more than one backup version exists, a user can restore the active backup version of the file or any inactive backup versions.

## Archive and Retrieve

To preserve files for later use or for records, a user can request ADSM to copy files, subdirectories, and directories for long-term storage on media controlled by ADSM. After users archive a file, they can save disk space by erasing the original file from their workstation.

When a user retrieves a file, ADSM sends a copy of the file to the client node. The archived file remains in ADSM storage.

## Migration and Recall

If the Hierarchical Storage Management (HSM) feature of ADSM is activated on a client node, users can migrate files from client node storage to data storage and recall files to the client node as needed. HSM frees space on client nodes for new data and makes more efficient use of your storage.

For details about using HSM on clients, see *ADSM Using the UNIX HSM Clients*.

### Migration

When a file is migrated to the server, it is replaced on the client node with a small stub file of the same name as the original file. The stub file contains data needed to locate the migrated file on server storage.

ADSM provides selective and automatic migration. Selective migration lets users migrate files by name. The two types of automatic migration are:

**Threshold** If space usage exceeds a high threshold set at the client node, migration begins and continues until usage drops to the low threshold also set at the client node.

**Demand** If an out-of-space condition occurs for a client node, migration begins and continues until usage drops to the low threshold.

To prepare for efficient automatic migration, ADSM copies a percentage of user files from the client node to the server. The *premigration* process occurs whenever ADSM completes an automatic migration. The next time free space is needed at the client node, the premigrated files at the server can be quickly changed to migrated files. The default premigration percentage is the difference between the high and low thresholds.

Files are selected for automatic migration and premigration based on the number of days since the file was last accessed and also on other factors set at the client node.

### Recall

ADSM provides selective and transparent recall. Selective recall lets users recall files by name. Transparent recall occurs automatically when a user accesses a migrated file.

### Reconciliation

Migration and premigration can create inconsistencies between client node and server storage. For example, if a user deletes a migrated file from the client node, the copy remains at the server. At regular intervals set at the client node, ADSM compares client node and server storage and reconciles the two by deleting from the server any outdated files or files that do not exist at the client node.

---

## Policy Objects

Policy administrators specify how files are backed up, archived, migrated from client node storage, and managed in ADSM storage. This process defines policy objects that are used to implement ADSM policies. The following figure shows the objects and their relationships:

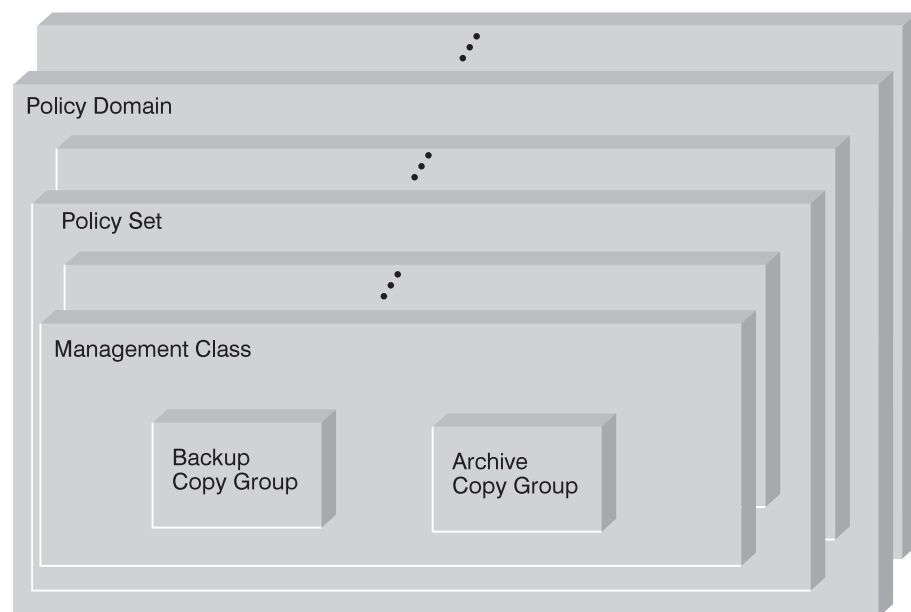


Figure 26. ADSM Policy Objects

### **Backup copy group**

Controls how ADSM performs backup processing of files associated with it. A backup copy group determines the following:

- If a file is backed up (even if it has not changed since the last backup)
- How many days must elapse before a file can be backed up again
- How to handle files that are in use during backup
- Where the server stores backup versions of files and directories
- How many backup versions the server keeps of files and directories
- How long the server keeps backup versions of files and directories

### **Archive copy group**

Controls how ADSM performs archive processing of files associated with it. An archive copy group determines the following:

- How to handle files that are in use during archive
- Where the server stores archived copies of files
- How long the server keeps archived copies of files

### **Management class**

Associates backup and archive groups with files and specifies if and how client node files are migrated to storage pools. A management class can contain one backup copy group, one archive copy group, both a backup and archive copy group, or no copy groups. Users can *bind* (that is, associate) their files to a management class.

### **Policy set**

Specifies the management classes that are available to groups of users. Policy sets contain one or more management classes: a *default management class* and any number of additional management classes.

### **Policy domain**

Lets an administrator group client nodes by the policies that govern their files and by the administrators who manage their policies. A policy domain contains one or more policy sets, but only one policy set (named ACTIVE) can be active at a time. ADSM uses the active policy set to manage files for client nodes assigned to a policy domain.

You can use policy domains to:

- Provide default storage management policies
- Group client nodes with similar storage management requirements
- Restrict the number of management classes to which users have access

---

## **Management Classes**

Each user is assigned to a single policy domain, and the user has access only to the management classes contained in the domain. The management classes specify whether or not space management is to be done. The copy groups in these management classes specify the number of backup versions retained in ADSM storage and the length of time to retain backup versions and archive copies.

For example, if most users at an installation need only one backup version of a file. You can create a policy domain that contains only one management class whose backup copy group allows only one backup version. Then you can assign these



client nodes to the policy domain. See “Registering Client Nodes” on page 104 for information on registering client nodes and assigning policy domains to them.

## Management Class Configuration

Before defining a management class, consider whether the management class should contain:

### A backup copy group and an archive copy group

For example, most users need to back up and archive documents, spread sheets, and graphics.

### A backup copy group only

For example, some users only want to back up application files (such as database, log, or history files that change daily).

### An archive copy group only

A management class that contains only an archive copy group is useful for users who create:

- Point-in-time files. For example, an engineer can archive the design of an electronic component and the software that created the design. Later, the engineer can use the design as a base for a new electronic component.
- Files that are rarely used but need to be retained for a long time. A user can erase the original file without affecting how long the archive copy is retained in data storage. Examples include legal records, patient records, and tax forms.

### Neither a backup nor an archive copy group

A management class that contains neither a backup nor an archive copy group prevents a file from ever being backed up or archived. If users bind their files to a management class without copy groups, ADSM issues warning messages.

**Warning:** This type of management class is not recommended for most users. Use such a management class carefully to prevent users from mistakenly selecting it.

## Default Management Classes

Each policy set must include a default management class, which is used:

- To manage files that are not bound to a specific management class, as defined by the INCLUDE option in the include-exclude list.
- To manage existing backup versions when a management class name is deleted from the server as described in “How Files Are Associated with a Management Class” on page 118.
- To manage existing archive copies when a management class is deleted from the server. ADSM does not rebind archive copies but does use the archive copy group (if one exists) in the default management class.

A typical default management class should do the following:

- Meet the storage management needs for most of your users
- Contain both a backup copy group and an archive copy group
- Set serialization static or shared static to ensure the integrity of backed up and archived files

- Retain backup versions and archive copies for a sufficient amount of time
- Retain directories for at least as long as any files are associated with the directory

Other management classes can contain copy groups tailored either for the needs of special sets of users or for the needs of most users under special circumstances.

## The Include-Exclude List

A user can define an include-exclude list to specify which files are eligible for backup services, which files can be migrated, and how ADSM manages backed up, archived, and migrated files.

If a user does not create an include-exclude list:

- All files belonging to the user are eligible for backup services.
- The default management class governs backup, archive, and migration.

With an include-exclude list, users can:

- Exclude files or directories from backup and client migration operations

For example, the following figure shows that the SSTEINER node ID excludes all core files from being eligible for backup and client migration.

- Include any previously excluded files

For example, the following figure shows that the SSTEINER node ID includes the /home/ssteiner/options.scr file. This file is now eligible for backup and client migration.

- Bind a file to a specific management class

For example, the following figure shows that all files and subdirectories belonging to the /home/ssteiner/driver5 directory are managed by the criteria defined in the MCENGBK2 management class.

```
exclude ../../core
include /home/ssteiner/options.scr
include /home/ssteiner/driver5/../../* mcengbk2
```

For information on how to create an include-exclude list, see the user's publication for the appropriate client.

## How Files Are Associated with a Management Class

*Binding* is the process of associating files with a management class. The policies defined in the management class then apply to the bound files. A user binds a file to a management class name by using:

- The INCLUDE option in an include-exclude list
- The ARCHMC option when archiving a file
- The DIRMC option when backing up a file

See the user's publication for the appropriate client for details.

When a user does not bind a file to a management class, the client node binds the file to the default management class in the active policy set.

A file remains bound to a management class name even if the attributes of the management class change. The following scenario illustrates this process:

1. A file named REPORT.TXT is bound to the default management class that contains a backup copy group specifying that up to three backup versions can be retained in data storage.
2. During the next week, three backup versions of REPORT.TXT are stored in ADSM storage. The active and two inactive backup versions are bound to the default management class.
3. The administrator assigns a new default management class that contains a backup copy group specifying only up to two backup versions.
4. The administrator then activates the policy set, and the new default management class takes effect.
5. Expiration processing occurs (see “Expiration Processing” for details). REPORT.TXT is still bound to the default management class, which now includes new retention criteria. Therefore, the oldest inactive version is expired, and one active and one inactive backup version remain in storage.

*Rebinding* is the process of associating a file with a new management class. Backup versions of files are rebound in the following cases:

- The user changes the management class specified in the include-exclude list and does a backup.
- The user specifies a different management class by using the DIRMC option when doing a backup.
- An administrator activates a policy set that does not contain a management class with the same name.
- An administrator assigns a client node to a different policy domain, and the active policy set in that policy domain does not have a management class with the same name.

If a file is bound to a management class that no longer exists, ADSM uses the default management class to manage the backup versions. When the user does another backup, ADSM rebinds the file and any backup versions to the default management class.

**Note:** Archive copies are never rebound because each archive operation creates a different archive copy. Archive copies remain bound to the management class name specified when the user archived them.

## Expiration Processing

Backup and archive copy groups can specify the criteria that make copies of files eligible for deletion from data storage. However, even when a file becomes eligible for deletion, the file is not deleted until expiration processing occurs. If expiration processing does not occur periodically, storage pool space is not reclaimed from expired client files, and the ADSM server requires increased storage space.

See “Starting Expiration Processing” on page 135 for details about how to invoke expiration processing.

---

## File Eligibility for Policy Operations

This section describes how ADSM selects files for the following operations:

- Incremental backup
- Selective backup
- Archive
- Migration from a client node (hierarchical storage management)

### Incremental Backup

When a user requests an incremental backup, ADSM performs the following steps to determine eligibility:

1. Checks each file against the user's include-exclude list:
  - Files that are excluded are not eligible for backup.
  - If files are not excluded and a management class is specified with the INCLUDE option, ADSM uses that management class.
  - If files are not excluded but a management class is not specified with the INCLUDE option, ADSM uses the default management class.
  - If no include-exclude list exists, all files in the client domain are eligible for backup, and ADSM uses the default management class.
2. Checks the management class of each included file:
  - If there is a backup copy group, ADSM goes to step 3.
  - If there is no backup copy group, the file is not eligible for backup.
3. Checks the *mode*, *frequency*, and *serialization* defined in the backup copy group.

**Mode** Specifies if the file is backed up only if it has changed since the last backup (*modified*) or whenever a backup is requested (*absolute*).

**Frequency** Specifies the minimum number of days that must elapse between backups and how files are handled if they are modified during backup processing.

**Serialization** Specifies how files are handled if they are modified while being backed up and what ADSM does if modification occurs.

- If the mode is modified and the minimum number of days have elapsed since the file was last backed up, the server determines if the file has been changed since it was last backed up:
  - If the file has been changed and the serialization requirement is met, the file is backed up.
  - If the file has not been changed, it is not backed up.
- If the mode is modified and the minimum number of days have not elapsed, the file is not eligible for backup.
- If the mode is absolute, the minimum number of days have elapsed since the file was last backed up, and the serialization requirement is met, the file is backed up.
- If the mode is absolute and the minimum number of days have not elapsed, the file is not eligible for backup.

## Selective Backup

When a user requests a selective backup, ADSM performs the following steps to determine eligibility:

1. Checks the file against any include or exclude statements contained in the user include-exclude list:
  - Files that are not excluded are eligible for backup. If a management class is specified with the INCLUDE option, ADSM uses that management class.
  - If no include-exclude list exists, only files specified on the command line are eligible for backup, and ADSM uses the default management class.
2. Checks the management class of each included file:
  - If the management class contains a backup copy group and the serialization requirement is met, the file is eligible for backup. Serialization specifies how files are handled if they are modified while being backed up and what ADSM does if modification occurs.
  - If the management class does not contain a backup copy group, the file is not eligible for backup.

## Archive

When a user requests the archiving of a file or a group of files, ADSM performs the following steps to determine eligibility:

1. Checks the files against the user's include-exclude list to see if any management classes are specified:
  - ADSM uses the default management class for files that are not bound to a management class.
  - If no include-exclude list exists, ADSM uses the default management class unless the user specifies another management class. See the user's publication for the appropriate client for details.
2. Checks the management class for each file to be archived.
  - If the management class contains an archive copy group and the serialization requirement is met, the file is archived. Serialization specifies how files are handled if they are modified while being archived and what ADSM does if modification occurs.
  - If the management class does not contain an archive copy group, the file is not archived.

## Automatic Migration from a Client Node

A file is eligible for automatic migration from a client node if it meets all of the following criteria:

- It resides on a node on which the root user has added and activated hierarchical storage management.
- It is not excluded from migration in the include-exclude list.
- It meets management class requirements for migration:
  - The file is not a character special file, a block special file, a FIFO special file (that is, a named pipe file) or a directory.

- The file is assigned to a management class that calls for space management.
- The management class calls for automatic migration after a specified number of days, and that time has elapsed.
- A backup version of the file exists if the management class requires it.
- The file is larger than the stub file that would replace it (plus one byte) or the file system block size, whichever is larger.

End of Diagnosis, Modification or Tuning Information

---

## How Client Migration Works with Backup and Archive

As an administrator, you can define a management class that specifies automatic migration under certain conditions. For example, if the file has not been accessed for at least 30 days and a backup version exists, the file is migrated. You can also define a management class that allows users to selectively migrate whether or not a backup version exists. Users can also choose to archive files that have been migrated:

- If the file is backed up or archived to the server to which it was migrated, ADSM copies the file from the migration storage pool to the backup or archive storage pool. For a tape-to-tape operation, each storage pool must have a tape drive.
- If the file is backed up or archived to a different server, ADSM accesses the file by using the migrate-on-close recall mode.

The file resides on the client node only until ADSM stores the backup version or the archived copy in the backup or archive storage pool.

When users restore a backup version of a migrated file, ADSM deletes the migrated copy of the file from server storage the next time reconciliation is run.

If users do not specify that the file is to be erased from the client after it is archived, the file remains migrated. If users specify that the file is to be erased, ADSM deletes the migrated file from ADSM storage the next time reconciliation is run.

The default management class delivered with ADSM specifies that a backup version of a file must exist before the file is eligible for migration.

---

## Using the Standard Storage Management Policies

ADSM provides a set of policy objects, named STANDARD. If you use these standard objects, you can begin using ADSM immediately.

When you register a client node, the default is to assign the node to the STANDARD policy domain. If users register their own workstations during open registration, they are also assigned to the STANDARD policy domain.

ADSM provides a standard policy domain, policy set, management class, backup copy group, and archive copy group. Each policy object is named STANDARD. The following table shows the attributes of the ADSM-supplied objects.

<i>Table 13. Standard Policy Objects Provided with ADSM</i>	
<b>Standard Policy Domain</b>	
<ul style="list-style-type: none"> <li>• Backup retention grace period = 30 days</li> <li>• Archive retention grace period = 365 days</li> </ul>	
<b>Standard Policy Set (ACTIVE)</b>	
<ul style="list-style-type: none"> <li>• Default Management class = STANDARD</li> </ul>	
<b>Standard Management Class</b>	
<ul style="list-style-type: none"> <li>• No space management</li> </ul>	
<b>Standard Backup Copy Group</b>	<b>Standard Archive Copy Group</b>
<ul style="list-style-type: none"> <li>• Destination = BACKUPPOOL</li> <li>• Backup frequency = 0 days</li> <li>• Incremental backup of modified files</li> <li>• Files cannot be backed up if in use</li> <li>• Retain up to 2 versions</li> <li>• Retain version for 60 days</li> </ul>	<ul style="list-style-type: none"> <li>• Destination = ARCHIVEPOOL</li> <li>• Files cannot be archived if in use</li> <li>• Retain archive copy for 365 days</li> </ul>

## Creating Your Own Storage Management Policies

Task	Required Privilege Class
Define or copy a policy domain	System
Update a policy domain over which you have authority	Restricted policy
Define, update, or copy policy sets and management classes in any policy domain	System or unrestricted policy
Define, update, or copy policy sets and management classes in policy domains over which you have authority	Restricted policy
Define or update copy groups in any policy domain	System or unrestricted policy
Define or update copy groups that belong to policy domains over which you have authority	Restricted policy
Assign a default management class to a nonactive policy set in any policy domain	System or unrestricted policy
Assign a default management class to a nonactive policy set in policy domains over which you have authority	Restricted policy
Validate and activate policy sets in any policy domain	System or unrestricted policy
Validate and activate policy sets in policy domains over which you have authority	Restricted policy
Start inventory expiration processing	System

You may need more flexibility in your storage management policies than the standard ADSM policy objects provide. If so, you can create your own policies in either of two ways: you can define the objects by specifying each attribute, or you can copy existing objects and update only those attributes that you want to change. The following table shows another advantage of copying objects: some associated objects are copied in a single operation.

If you copy:	You create:
Policy Domain	A new policy domain with: <ul style="list-style-type: none"> <li>• A copy of each policy set from the original domain</li> <li>• A copy of each management class in each original policy set</li> <li>• A copy of each copy group in each original management class</li> </ul>
Policy Set	A new policy set <b>in the same policy domain</b> with: <ul style="list-style-type: none"> <li>• A copy of each management class in the original policy set</li> <li>• A copy of each copy group in the original management class</li> </ul>
Management Class	A new management class <b>in the same policy set</b> and a copy of each copy group in the management class

The rest of this chapter describes the tasks involved in creating new storage management policies for your installation:

1. Define policy domains to manage groups of client nodes. See page 126.
2. Define policy sets for different storage management policies. See page 127.
3. Define management classes to match users' storage management requirements. See page 127.
4. Define backup copy groups to specify which files can be backed up and how to manage backup versions. See page 128.
5. Define archive copy groups to specify whether a file can be archived if it is in use and to manage archive copies. See page 132.
6. Assign a default management class to each policy set to match the most common storage management requirements of client nodes in the policy domain. See page 134.
7. Validate all policy sets, and activate one policy set for each policy domain. See page 135.
8. Start expiration processing. See page 135.

To help users take advantage of ADSM, you can set up the policy environment by doing the following:

- Create include-exclude lists for inexperienced users or for users who have simple storage management needs
- Provide a sample include-exclude list to users who want to specify how ADSM manages their files. You can show users who prefer to manage their own files how to:
  - Request information about management classes.
  - Select a management class that meets backup and archive requirements.
  - Use include-exclude lists to bind management classes to their files.

For information on how to create an include-exclude list, see the user's publication for the appropriate client.

- Automate incremental back up procedures by defining schedules for each policy domain. Then associate schedules with client nodes in each policy domain. For information on schedules, see Chapter 7, "Scheduling Operations" on page 141.



## Example: Sample Policy Objects

The following figure shows the policies for an engineering department. This example is used throughout the rest of this chapter.

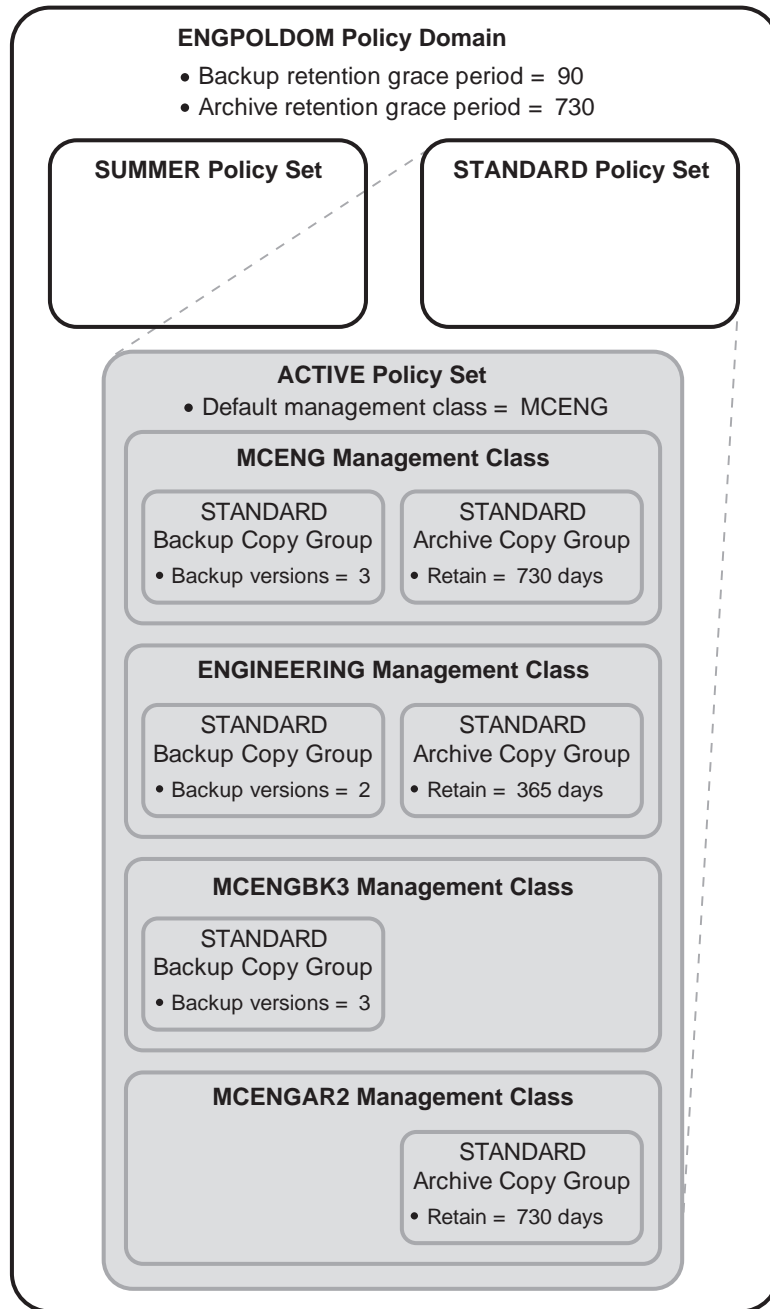


Figure 27. An Example of Policy Objects Defined for an Engineering Department

The domain contains two policy sets, **STANDARD** and **SUMMER**. The policy set named **STANDARD** is active. Only one policy set can be active at a time. When a policy set is activated, the server makes a copy of the policy set and names it **ACTIVE**.

The ACTIVE policy set contains four management classes: ENGINEERING, MCENG, MCENGBK3, and MCENGAR2. The default management class is MCENG.

## Defining and Updating a Policy Domain

When you update or define a policy domain, you specify:

### Backup Retention Grace Period

Specifies the number of days to retain a backup version when the server cannot rebind the file to an appropriate management class. The backup retention grace period protects backup versions from being immediately deleted when:

- A user binds a file to a new management class name that does not contain a backup copy group
- The management class to which a file is bound no longer exists and the default management class does not contain a backup copy group

When users try to back up the file and the management class does not contain a backup copy group, ADSM uses the backup retention grace period to manage all existing backup versions of the file, and the file is not backed up.

Backup versions of the file are retained in data storage only for the backup retention grace period. This period starts from the day of the backup, unless a user binds the file to a management class containing a backup copy group and does a backup.

For example, if the backup retention grace period for the STANDARD policy domain is used and set to 30 days, backup versions using the grace period expire in 30 days from the day of the backup.

### Archive Retention Grace Period

Specifies the number of days to retain an archive copy when the server cannot rebind the file to an appropriate management class. The retention grace period protects archive copies from being immediately deleted if the default management class does not contain an archive copy group.

The archive copy of the file is retained in data storage for the number of days specified by the archive retention grace period starting from the day on which the file is first archived.

For example, if the archive retention grace period for the policy domain STANDARD is used, an archive copy expires 365 days from the day the file is first archived.

## Example: Defining a Policy Domain

To create a new policy domain you can do one of the following:

- Copy an existing policy domain and update the new domain
- Define a new policy domain from the beginning

**Note:** When you copy an existing domain, you also copy any associated policy sets, management classes, and copy groups.

For example, to copy and update, follow this procedure:

1. Copy the STANDARD policy domain to the ENGPOLDOM policy domain by entering:

```
copy domain standard engpoldom
```

ENGPOLDOM now contains the standard policy set, management class, backup copy group, and archive copy group.

2. Update the policy domain ENGPOLDOM so that the backup retention grace period is extended to 90 days and the archive retention grace period is extended to 2 years by entering:

```
update domain engpoldom description='Engineering Policy Domain' -  
backretention=90 archretention=730
```

## Defining and Updating a Policy Set

When you define or update a policy set, specify:

### Policy domain name

Names the policy domain to which the policy set belongs

### Example: Defining a Policy Set

A business with seasonal employees needs two policy sets. During most of the year, most users would use the STANDARD policy set. During the summer, it would activate the SUMMER policy set to provide new management classes for users who are seasonal employees. To create the SUMMER policy set in the STANDARD policy domain, the business would perform the following steps:

1. Copy the STANDARD policy set and name the new policy set SUMMER:

```
copy policyset standard standard summer
```

**Note:** When you copy an existing policy set, you also copy any associated management classes and copy groups.

2. Update the description of the policy set named SUMMER, enter:

```
update policyset standard summer -  
description='Policy set activated during summer for STANDARD domain'
```

## Defining and Updating a Management Class

When you define or update a management class, specify:

### Policy domain name

Names the policy domain to which the management class belongs.

### Policy set name

Names the policy set to which the management class is assigned.

**Whether space management is to be done**

Specifies that the files are eligible for both automatic and selective migration, only selective migration, or no migration.

**How frequently files can be migrated**

Specifies the minimum number of days that must elapse since a file was last accessed before it is eligible for automatic migration.

**Whether backup is required**

Specifies whether a backup version of a file must exist before the file can be migrated.

**Where the files are to be stored**

Specifies the name of the storage pool in which migrated files are stored. Your choice could depend on factors such as:

- The number of client nodes migrating to the storage pool. When many user files are stored in the same storage pool, volume contention can occur as users try to migrate files to or recall files from the storage pool.
- How quickly the files must be recalled. If users need immediate access to migrated versions, you can specify a disk storage pool as the destination.

**Note:** You cannot specify a copy storage pool as a destination.

**Example: Define a New Management Class**

Create a new management class containing a backup copy group and an archive copy group:

1. Copy the STANDARD management class from the STANDARD policy set to the new management class (named MCENG) by entering:

```
copy mgmtclass engpoldom standard standard mceng
```

The server copies the management class description, standard backup copy group, and standard archive copy group to MCENG.

2. Update the description of the MCENG management class by entering:

```
update mgmtclass engpoldom standard mceng -  
description='Engineering Mgmt Class with Backup & Archive Copy Groups'
```

## Defining and Updating a Backup Copy Group

To define or update a backup copy group on the graphical user interface or command line, specify:

**Where files are to be stored**

Specifies a defined storage pool. Your choice can depend on factors such as:

- The number of client nodes backing up to the storage pool. When many user files are stored in the same storage pool, volume contention can occur as users try to back up to or restore files from the storage pool.

- How quickly the files must be restored. If users need immediate access to backup versions, you could specify a disk storage pool as the destination.

**Note:** You cannot specify a copy storage pool.

### **If files can be modified during backup**

Specifies how files are handled if they are modified while being backed up and what ADSM does if modification occurs. This attribute, called serialization, can be one of four values:

#### **Static**

Specifies that if the file or directory is modified during a backup, ADSM does not back it up. ADSM does not retry the backup.

#### **Shared Static**

Specifies that if the file or directory is modified during an backup, ADSM does not back it up. However, ADSM retries the backup as many times as specified by the CHANGINGRETRIES option in the client options file.

#### **Dynamic**

Specifies that a file or directory is backed up on the first attempt, even if the file or directory is being modified during the backup.

#### **Shared Dynamic**

Specifies that if a file or directory is modified during a backup attempt, ADSM backs it up on its last try even if the file or directory is being modified. ADSM retries the backup as many times as specified by the CHANGINGRETRIES option in the client options file.

For most files, set serialization to either static or shared static to prevent the server from backing up a file while it is being modified.

**Warning:** If a file is backed up while it is in use (shared dynamic or dynamic serialization), the copy may not contain all the changes and may not be usable.

However, you may want to define a copy group with a serialization of shared dynamic or dynamic for files where log records are continuously added, such as an error log. If you only have copy groups that use static or shared static, these files may never be backed up because they are constantly in use. With shared dynamic or dynamic, the log files are backed up. However, the backup version may contain a truncated message.

### **How frequently files can be backed up**

Specifies the minimum number of days that must elapse between incremental backups. Frequency works with the mode parameter, which specifies whether a file or directory is considered for incremental backup only if it has changed since the last backup or regardless of whether it has been changed. ADSM does not check this attribute when a user requests a selective backup for a file. You can select from two modes:

#### **Modified**

A file is considered for incremental backup only if it has changed since the last backup. A file is considered changed if any of the following items is different:

- Date on which the file was last modified
- File size
- File owner

- File permissions

### **Absolute**

A file is considered for incremental backup regardless of whether it has changed since the last backup.

For example, if frequency is 3 and mode is modified, a file or directory is backed up only if it has been changed and if three days have passed. If frequency is 3 and mode is absolute, a file or directory is backed up after three days have passed whether or not the file has changed.

Use the modified mode when users want to retain multiple backup versions. If the mode is set to absolute, users may have three *identical* backup versions, rather than three different backup versions.

Absolute mode can be useful for forcing a full backup or ensuring that OS/2 files with extended attributes are backed up because ADSM does not detect changes to the extended attributes.

When you set the mode to absolute, set frequency to 0 if you want to ensure that a file is considered for backup each time incremental backups are scheduled for or initiated by a user.

### **How many backup versions to retain**

Specifies the number of backup versions. Multiple versions of files are useful when users continually update files and sometimes need to restore the original file from which they started. Two parameters determine how many active and inactive backup copies to retain:

#### **Versions Data Exists**

The maximum number of different backup versions that the server retains for files and directories currently on the workstation.

If users select a management class that allows more than one backup version, the most current version is called the *active* version. All other versions are called *inactive* versions.

For example, in Figure 28 on page 131, the most current version of REPORT.TXT was created on Friday at 3 p.m. There are two inactive versions of REPORT.TXT.

When the maximum number of backup versions is exceeded, the server deletes the oldest version.

For example, if the maximum number of versions allowed for MEMO.DAT is 3, and a user runs a backup process that creates a fourth version, the server deletes the oldest version. In this example, the backup version created on Thursday at 8:05 a.m. is deleted from data storage.

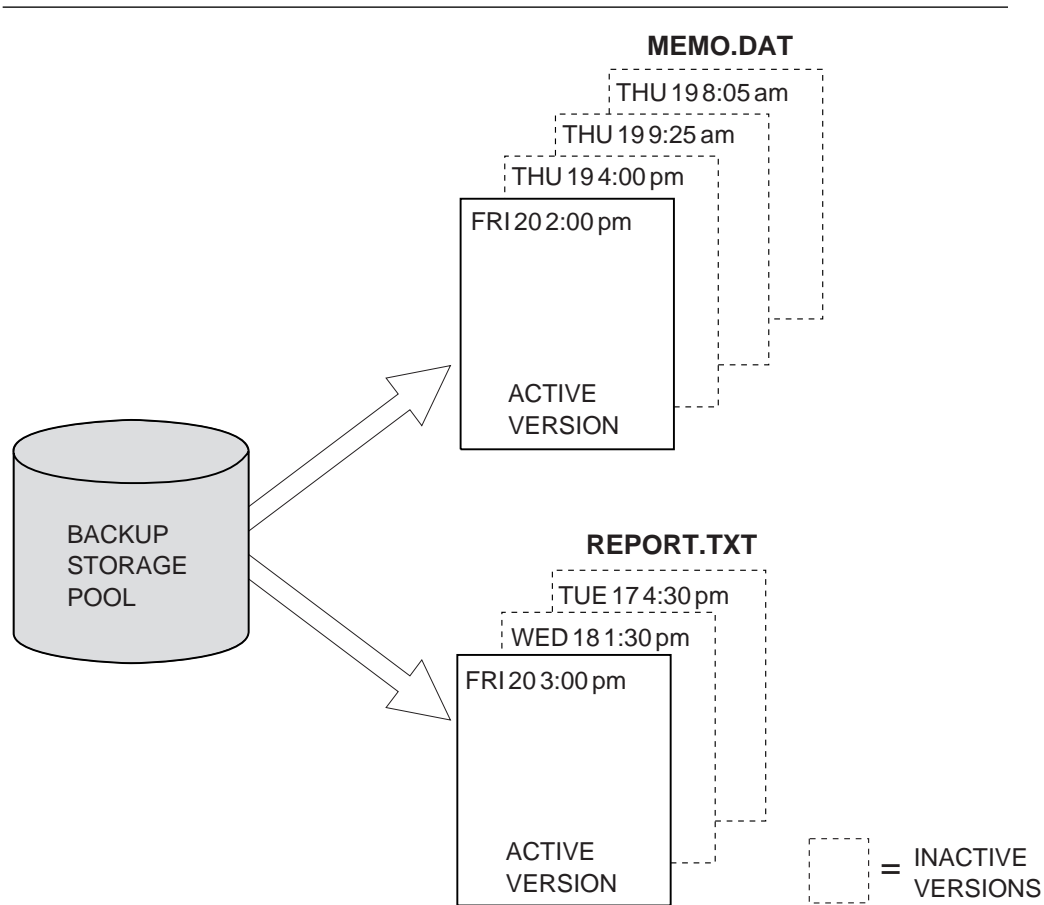


Figure 28. Example of Active and Inactive Versions of Backed Up Files

### Versions Data Deleted

The maximum number of different backup versions that the server retains for files and directories that have been erased from a workstation. The server ignores this parameter while the file or directory remains on the workstation.

If users erase a file or directory from their client node, then the next time a backup process is run, the server changes the active backup version to inactive and erases the oldest versions that are more than the number specified by this parameter.

The expiration date for the remaining versions is based on the Retain Extra Versions and Retain Only Version parameters.

### How long to retain files in storage

Specifies two parameters that determine specify how long to retain backup versions:

#### Retain Extra Versions

Specifies the retention time, in days, for all but the most recent backup version. The value of this parameter determines which versions are deleted during inventory expiration processing.

If NOLIMIT is specified, inactive backup versions are deleted based on the Versions Data Exists or Versions Data Deleted parameters.

### Retain Only Version

Specifies how many days ADSM retains the only backup version it has of a file when the original file has been deleted from the workstation.

If NOLIMIT is specified, the last version is retained forever unless a user or administrator deletes the file from data storage.

### Example: Define a Backup Copy Group

Define a backup copy group belonging to the MCENG management class in the STANDARD policy set belonging to the ENGPOLDOM policy domain. This new copy group must do the following:

- Let users back up changed files, regardless of how much time has elapsed since the last backup
- Retain up to 4 inactive backup versions when the original file resides on the user workstation
- Retain up to 3 inactive backup versions when the original file is deleted from the user workstation
- Retain extra inactive backup versions for 90 days
- If there is only one backup version, retain it for 600 days
- Prevent files from being backed up if they are in use
- Store files in the ENGBACK1 storage pool

To define the backup copy group, enter:

```
define copygroup engpoldom standard mceng standard -  
destination=engback1 serialization=static -  
verexists=5 verdeleted=4 retextra=90 retonly=600
```

## Defining and Updating an Archive Copy Group

To define or update an archive copy group on the graphical user interface or command line, specify:

### Where files are to be stored

Specifies a a defined storage pool. Your choice can depend on factors such as:

- The number of client nodes backing up to the storage pool. When many user files are stored in the same storage pool, volume contention can occur as users archive files to and retrieve files from the storage pool.
- How quickly the files must be restored. If users need immediate access to archive copies, you could specify a disk storage pool as the destination.

**Note:** You cannot specify a copy storage pool as a destination.

### If files can be modified during archive

Specifies how files are handled if they are modified while being archived and what ADSM does if modification occurs. There are four options that you can chose from: This attribute, called serialization, can be one of four values:



**Static**

Specifies that if the file or directory is modified during an archiving process, ADSM does not archive it. ADSM does not retry the archive.

**Shared Static**

Specifies that if the file or directory is modified during an archive process, ADSM does not archive it. However, ADSM retries the archive process as many times as specified by the CHANGINGRETRIES option in the client options file.

**Dynamic**

Specifies that a file or directory is archived on the first attempt, even if the file or directory is being modified during the archive process.

**Shared Dynamic**

Specifies that if a file or directory is modified during the archive attempt, ADSM archives it on its last try even if the file or directory is being modified. ADSM retries the archive process as many times as specified by the CHANGINGRETRIES option in the client options file.

For most files, set serialization to either static or shared static to prevent the server from archiving a file while it is being modified.

**Warning:** If a file is archived while it is in use (shared dynamic or dynamic serialization), the copy may not contain all the changes and may not be usable.

However, you may want to define a copy group with a serialization of shared dynamic or dynamic for files where log records are continuously added, such as an error log. If you only have copy groups that use static or shared static, these files may never be archived because they are constantly in use. With shared dynamic or dynamic, the log files are archived. However, the archive copy may contain a truncated message.

**How long to retain an archived copy**

Specifies the number of days to retain an archived copy in storage. When the time elapses, ADSM deletes the file.

**Example: Define an Archive Copy Group**

Define an archive copy group belonging to the MCENG class that:

- Allows users to archive a file if it is not in use
- Retains the archive copy for 730 days
- Stores files in the ENGARCH1 storage pool

To define a STANDARD archive copy group to the MCENG management class in the STANDARD policy set belonging to the ENGPOLDOM policy domain, enter:

```
define copygroup engpoldom standard mceng standard -  
type=archive destination=engarch1 serialization=static -  
retver=730
```

## Assigning a Default Management Class

After you have defined your policy sets and the management classes that they contain, you must assign a default management class for each policy set. See “Default Management Classes” on page 117 for suggestions about the content of default management classes.

### Example: Assign a Default Management Class

To assign the STANDARD management class as the default management class for the SUMMER policy set in the STANDARD policy domain, enter:

```
assign defmgmtclass standard summer standard
```

The default management class is copied from the STANDARD policy set to the SUMMER policy set. Before the new default management class takes effect, you must activate the policy set.

## Validating and Activating Policy Sets

After you have defined your policy sets and assigned management classes to them, you can validate those policy sets and activate one policy set for the policy domain.

### Validating Policy Sets

When you validate a policy set, the server examines the management class and copy group definitions in the specified policy set and reports on conditions that need to be considered if the policy set is activated.

Validation fails if the policy set does not contain a default management class. The following conditions result in warning messages during validation:

- The storage destinations specified for backup, archive, or migration do not refer to defined storage pools.

A backup, archive, or migration operation fails when the operation involves storing a file in a nonexistent storage pool.

- A storage destination specified for backup, archive, or migration is a copy storage pool.
- The default management class does not contain a backup or archive copy group.

When the default management class does not contain a backup or archive copy group, any user files bound to the default management class *are not* backed up or archived.

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

When users back up files that were bound to a management class that no longer exists in the active policy set, backup versions are rebound to the default management class. See “How Files Are Associated with a Management Class” on page 118 for details.

When the management class to which an archive copy is bound no longer exists and the default management class does not contain an archive copy

group, the archive retention grace period is used to retain the archive copy. See “Defining and Updating a Policy Domain” on page 126 for details.

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

When users perform a backup and the backup copy group no longer exists in the management class to which a file is bound, backup versions are managed by the backup retention grace period, and the workstation file is not backed up. See “Defining and Updating a Policy Domain” on page 126.

- 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.

## Activating Policy Sets

To activate a policy set, specify a policy domain and policy set name. When you activate a policy set, the server:

- Performs a final validation of the contents of the policy set
- Copies the original policy set to the active policy set

After a policy set has been activated, the original and the ACTIVE policy sets are two separate objects. For example, updating the original policy set has no effect on the ACTIVE policy set. You cannot update the ACTIVE policy set. To change its contents, you must do the following:

1. Copy the ACTIVE policy set to a policy set with another name.
2. Update the new policy set.
3. Validate the new policy set.
4. Activate the new policy set to have the server use the changes.

### Example: Validating and Activating a Policy Set

Validating and activating the SUMMER policy set in the STANDARD policy domain is a two-step process:

1. To validate the SUMMER policy set, enter:

```
validate policysset standard summer
```

2. To activate the SUMMER policy set, enter:

```
activate policysset standard summer
```

## Starting Expiration Processing

Copies of files that have expired are not deleted from data storage until expiration processing occurs. You can invoke expiration processing either automatically or by command. Automatic expiration processing can be controlled by the EXPINTERVAL parameter specified in the ADSM options file (dsmserv.opt). For details, see *ADSM Installing the Server and Administrative Client*. You can manually start expiration processing by issuing the following command:

```
expire inventory
```

Expiration processing then deletes eligible backup versions and archive file copies.

## Querying Policy Objects

Task	Required Privilege Class
Query any policy domain, policy set, management class, or copy group	Any administrator

You can request information about the contents of ADSM policy objects. For example, you might want to do this before creating new objects or helping users to choose a policies that fit their needs.

You can specify the output of a query in either standard or detailed format. The examples in this book are in standard format. Refer to *ADSM Administrator's Reference* for examples of detailed format output.

## Querying Copy Groups

To request information about backup copy groups (the default) in the ENGPOLDOM engineering policy domain, enter:

```
query copygroup engpoldom * *
```

The following figure is the output from the query. It shows that the ACTIVE policy set contains two backup copy groups that belong to the MCENG and STANDARD management classes.

Policy Domain Name	Policy Set Name	Mgmt Class Name	Copy Group Name	Retain Version
ENGPOLDOM	ACTIVE	MCENG	STANDARD	730
ENGPOLDOM	ACTIVE	STANDARD	STANDARD	365
ENGPOLDOM	STANDARD	MCENG	STANDARD	730
ENGPOLDOM	STANDARD	STANDARD	STANDARD	365
ENGPOLDOM	SUMMER	MCENG	STANDARD	730
ENGPOLDOM	SUMMER	STANDARD	STANDARD	365

To request information about archive copy groups in the ENGPOLDOM engineering policy domain, enter:

```
query copygroup engpoldom * type=archive
```

## Querying Management Classes

To request information about management classes in the ENGPOLDOM engineering policy domain, enter:

```
query mgmtclass engpoldom * *
```

The following figure is the output from the query. It shows that the ACTIVE policy set contains the MCENG and STANDARD management classes.

Policy Domain Name	Policy Set Name	Mgmt Class Name	Default Mgmt Class ?	Description
ENGPOLDOM	ACTIVE	MCENG	Yes	Engineering Management Class with Backup and Archive Copy Groups
ENGPOLDOM	ACTIVE	STANDARD	No	
ENGPOLDOM	STANDARD	MCENG	Yes	Engineering Management Class with Backup and Archive Copy Groups versions
ENGPOLDOM	STANDARD	STANDARD	No	
ENGPOLDOM	SUMMER	MCENG	Yes	Engineering Management Class with Backup and Archive Copy Groups versions
ENGPOLDOM	SUMMER	STANDARD	No	

## Querying Policy Sets

To query the system for information about policy sets in the ENGPOLDOM engineering policy domain, enter:

```
query policysset engpoldom *
```

The following figure is the output from the query. It shows an ACTIVE policy set and two inactive policy sets, STANDARD and SUMMER.

Policy Domain Name	Policy Set Name	Default Mgmt Class Name	Description
ENGPOLDOM	ACTIVE	MCENG	Policy Set Activated During Summer
ENGPOLDOM	STANDARD		
ENGPOLDOM	SUMMER	MCENG	Policy Set Activated During Summer

## Querying Policy Domains

To request information about a policy domain (for example, to determine if any client nodes are registered to that policy domain), enter:

```
query domain *
```

The following figure is the output from the query. It shows that both the ENGPOLDOM and STANDARD policy domains have client nodes assigned to them.

Policy Domain Name	Activated Policy Set	Activated Default Mgmt Class	Number of Registered Nodes	Description
ENGPOLDOM	SUMMER	ENGMC	3	Engineering Policy Domain
STANDARD	STANDARD	STANDARD	3	Installed default policy domain.

## Deleting Policy Objects

Task	Required Privilege Class
Delete policy domains	System
Delete any policy sets, management classes, or copy groups	System or unrestricted policy
Delete policy sets, management classes, or copy groups that belong to policy domains over which you have authority	Restricted policy

In the following sections, note that you cannot delete the ACTIVE policy set or objects in that policy set. Also note that when you delete an object, you also delete any objects belonging to it.

## Deleting Copy Groups

You can delete a backup or archive copy group that does not belong to a management class in the ACTIVE policy set.

To delete the backup and archive copy groups belonging to the MCENG and STANDARD management classes in the SUMMER policy set, enter:

```
delete copygroup engpoldom summer mceng type=backup
delete copygroup engpoldom summer standard type=backup
delete copygroup engpoldom summer mceng type=archive
delete copygroup engpoldom summer standard type=archive
```

## Deleting Management Classes

You can delete a management class that does not belong to the ACTIVE policy set. To determine if it belongs to the ACTIVE policy set, issue the QUERY MGMTCLASS command.

To delete the MCENG and STANDARD management classes from the SUMMER policy set, enter:

```
delete mgmtclass engpoldom summer mceng
delete mgmtclass engpoldom summer standard
```

**Note:** When you delete a management class from a policy set, the server deletes the management class and all copy groups that belong to the management class in the specified policy domain.

## Deleting Policy Sets

Authorized administrators can delete any policy set other than the ACTIVE policy set. To determine if it is active, issue the QUERY POLICYSET command. To delete the SUMMER policy set from the ENGPOLDOM engineering policy domain, enter:

```
delete policyset engpoldom summer
```

**Note:** When you delete a policy set, the server deletes all management classes and copy groups that belong to the policy set within the specified policy domain.

## Deleting Policy Domains

You delete a policy domain that has no client nodes registered to it. To determine if any client nodes are registered to a policy domain, issue the QUERY DOMAIN command.

To delete the STANDARD policy domain, perform the following steps:

1. Request a list of all client nodes assigned to the policy domain by entering:

```
query node * domain=standard
```

2. If client nodes are assigned to the policy domain, remove them in either of the following ways:
  - Assign each client to a new policy domain. For example, create a macro by entering:

```
update node htang domain=engpoldom
update node tomc domain=engpoldom
update node pease domain=engpoldom
```

If the active policy set in ENGPOLDOM does not have the same management class names as in the active policy set of the STANDARD policy domain, then backup versions of files may be bound to a different management class name, as described in “How Files Are Associated with a Management Class” on page 118.

- Delete each node from the STANDARD policy domain.

3. Delete the policy domain by entering:

```
delete domain standard
```

**Note:** When you delete a policy domain, the server deletes the policy domain and all policy sets (including the ACTIVE policy set), management classes, and copy groups that belong to the policy domain.



---

## Chapter 7. Scheduling Operations

ADSM includes a central scheduling component that allows the automatic processing of administrative commands and client operations. Administrative and client schedules consist of commands that are automatically processed during a specific time period.

Administrative commands can be scheduled for use in tuning server operations and to start functions that require significant server or system resources. Automating these operations allows the administrator to ensure that server resources are available when needed by clients.

Administrators can use central scheduling to automate client operations. Automating client operations frees users from having to perform them manually.

The sections listed in the following table begin at the indicated pages.

Section	Page
<b>Concepts:</b>	
Central scheduling	141
Scheduling modes	142
<b>Tasks:</b>	
Setting scheduling modes	143
Coordinating server and client node schedules	144
Managing schedules and associations	148
Managing scheduled events	155

Most tasks presented in this chapter can be performed using either the graphical user interface or the command-line interface. Table 6 on page 39 shows whether a task can be performed on the graphical user interface, the command-line interface, or both.

For information about issuing ADSM commands, refer to *ADSM Administrator's Reference*. For help performing a task on the graphical user interface, see the procedure described in "Using Online Help" on page 15.

---

### Central Scheduling

ADSM lets you define, update, copy, query, and delete administrative command and client operation schedules. These schedules consist of administrative commands or client operations that are processed during a specified time period when the schedule is activated.

Each administrative command and each client operation is called an event. Each scheduled event is tracked by the server and recorded in the database. Event records can be deleted from the database as space requires.

---

## Scheduling Modes

The central scheduler uses the *client-polling* and *server-prompted* scheduling modes.

### Client-Polling Scheduling Mode

You can use the client-polling scheduling mode with all communication methods. With this mode, a client node queries the server at prescribed time intervals to obtain a schedule. When the scheduled start time begins, the client node performs the scheduled operation and sends the results to the server. The client node then queries the server for its next scheduled operation.

This sequence is illustrated in Figure 29.

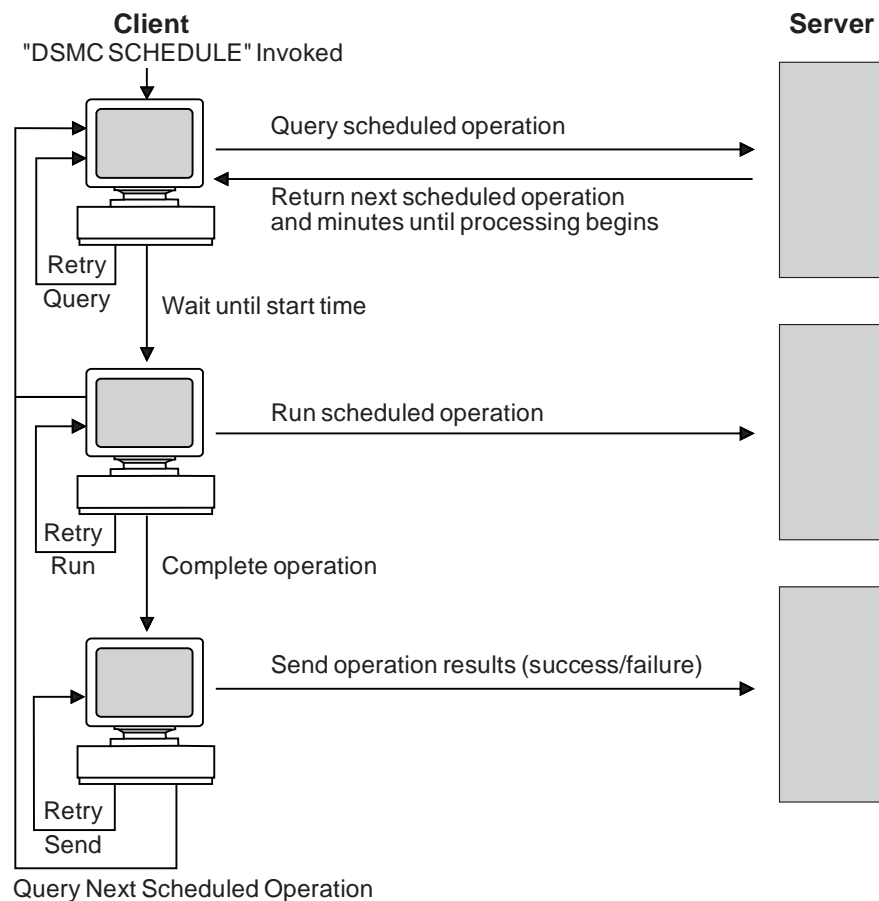


Figure 29. Client-Polling Scheduling Mode

### Server-Prompted Scheduling Mode

You can use the server-prompted scheduling mode only with client nodes that communicate with the server by using the TCP/IP communication method. With this mode, client nodes register their addresses with the server and then are contacted by the server when scheduled operations need to be performed and a session is available. When contacted, the client node queries the server for the operation, performs the operation, and then sends the results to the server.

This sequence is illustrated in Figure 30 on page 143.

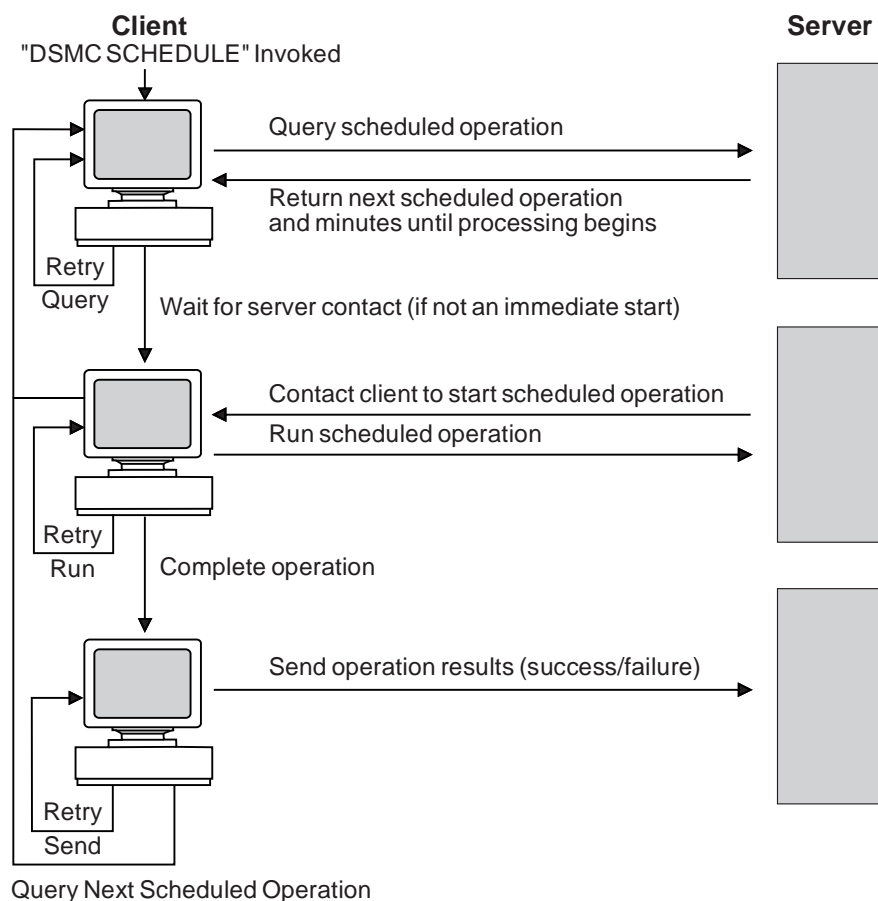


Figure 30. Server-Prompted Scheduling Mode

## Setting Scheduling Modes

Task	Required Privilege Class
Set scheduling modes	System

You can configure the server to use the client-polling scheduling mode, the server-prompted scheduling mode, or both modes.

Users can configure the client nodes to use either the client-polling scheduling mode or the server-prompted scheduling mode.

The following sections describe how to set the scheduling mode on the server and on client nodes.

## Setting the Scheduling Mode on the Server

You can set the scheduling mode on the server to client-polling, to server-prompted, or to both client-polling and server-prompted modes. The scheduling mode is set to both modes at installation.

If the scheduling mode is set to client-polling, the server cannot prompt client nodes:

```
set schedmodes polling
```

In this case, a client node with its scheduling mode set to server-prompted must wait until the server scheduling mode is set to server-prompted or to both server-prompted and client-polling for the scheduled work to begin.

If you set the scheduling mode to server-prompted, client nodes cannot poll the server:

```
set schedmodes prompted
```

In this case, a client node with its scheduling mode set to client-polling must wait until the server scheduling mode is set to client-polling or to both server-prompted and client-polling for the scheduled work to begin.

To let the server support both client-polling and server-prompted scheduling modes, set the scheduling mode to both modes:

```
set schedmodes any
```

In this case, any client node may set any scheduling mode and scheduled work will begin as specified.

## Setting the Scheduling Mode on Client Nodes

Users (root users on UNIX systems) set the scheduling mode on client nodes. They specify either the client-polling or the server-prompted scheduling mode on the command line or in the client user options file (client system options file on UNIX systems).

For more information, refer to the appropriate *ADSM Using the Backup-Archive Client*.

---

## Coordinating Server and Client Node Schedules

Task	Required Privilege Class
<ul style="list-style-type: none"> <li>• Set the maximum percentage of sessions for scheduled operations</li> <li>• Randomize schedule start times</li> <li>• Set how often clients query the server</li> <li>• Set the maximum number of times the client node scheduler can retry a command that failed</li> <li>• Set the time between retry attempts</li> </ul>	System

By coordinating server and client node schedules, can control the scheduler workload and client node contact with the server.

## Specifying the Schedule Period for Incremental Backup Operations

When you define a backup copy group, you specify the copy frequency, which is the minimum interval between successive backups. See “Defining and Updating a Backup Copy Group” on page 128. When you define a schedule, you specify the length of time between processings of the schedule. Do not process the schedule for incremental backups more often than the backup copy group frequency.

## Controlling the Scheduler Workload

An administrator can:

- Set the maximum percentage of concurrent client/server sessions for scheduled operations
- Randomize schedule start times

### Setting the Maximum Percentage of Sessions for Scheduled Operations

The number of concurrent client/server sessions is defined by the MAXSESSIONS option in the server options file, but you can set a maximum percentage of concurrent client/server sessions allowed for processing scheduled operations. Limiting the number of sessions available for scheduled operations ensures that sessions are available when users initiate any unscheduled operations, such as restoring or retrieving files, or backing up or archiving files.

If the number of sessions for scheduled operations is insufficient, you can increase either the total number of sessions or the maximum percentage of scheduled sessions. However, increasing the total number of sessions can adversely affect server performance, and increasing the maximum percentage of scheduled sessions can reduce the server opportunity to process unscheduled operations.

For example, assume that the maximum number of sessions between client nodes and the server is 80. If you want 25 percent of these sessions to be used by central scheduling, enter:

```
set maxschedsessions 25
```

The server allows 20 sessions to be used for scheduled operations.

For information about the MAXSESSIONS option, refer to *ADSM Installing the Server and Administrative Client*.

## Randomizing Schedule Start Times

To randomize a schedule start time means to scatter each schedule's start time across its startup window. A startup window is the start time and duration during which a schedule must be initiated.

The settings for randomization and the maximum percentage of scheduled sessions can affect whether schedules are successfully completed for client nodes. Users receive a message if all sessions are in use when they attempt to process a schedule. If this happens, you can increase randomization and the percentage of scheduled sessions allowed to make sure the server can handle the workload.

Increasing the size of the startup window (by increasing the schedule's duration) can also affect whether a schedule completes successfully. A larger startup window gives the client node more time to attempt initiation of a session with the server.

You might have to use trial and error to control the workload. To estimate how long client operations take, test schedules on several representative client nodes. Keep in mind, for example, that the first incremental backup for a client node takes longer than subsequent incremental backups.

For the client-polling scheduling mode, you can specify the percentage of the startup window that the server can use to randomize start times for different client nodes associated with a schedule.

If you set randomization to 0, no randomization occurs. This process can result in communication errors if many client nodes try to contact the server at the same instant.

The maximum percentage of randomization allowed is 50 percent. This limit ensures that half of the startup window is available for retrying scheduled commands that have failed.

It is possible, especially after a client node or the server has been restarted, that a client node may not poll the server until *after* the beginning of the startup window in which the next scheduled event is to start. In this case, the starting time is randomized over the specified percentage of the *remaining* duration of the startup window.

Consider the following situation:

- The startup window for a particular event is from 8:00 to 9:00
- Ten client nodes are associated with the schedule
- Nine client nodes poll the server before 8:00
- One client node does not poll the server until 8:30

To set randomization to 50 percent enter:

```
set randomize 50
```

The result is that the nine client nodes that polled the server *before* the beginning of the startup window are assigned randomly selected starting times between 8:00

and 8:30. The client node that polled at 8:30 receives a randomly selected starting time that is between 8:30 and 8:45.

## Controlling Contact with the Server

To control how often client nodes contact the server to perform a scheduled operation, an administrator can set:

- How often clients query the server
- The number of command retry attempts
- The amount of time between retry attempts

Users (root users on UNIX systems) can also set these values in their client user options files (client system options files for UNIX systems). However, user values are overridden by the values that the administrator specifies.

The client node communication paths to the server can vary widely with regard to response time or the number of gateways. In such cases, you can choose *not* to set these values so that users can tailor them for their own needs.

### Setting How Often Clients Query the Server

For the client-polling scheduling mode, you can specify the maximum number of hours the scheduler on a client node waits between attempts to contact the server to obtain a schedule.

You can set this period to correspond to the frequency with which the schedule changes are being made. If client nodes poll more frequently for schedules, changes to scheduling information (through administrator commands) are propagated more quickly to client nodes. However, increased polling by client nodes also increases network traffic.

If you want to have all clients using polling mode contact the server every 24 hours, enter:

```
set querieschedperiod 24
```

### Setting the Number of Command Retry Attempts

You can specify the maximum number of times the scheduler on a client node can retry a scheduled command that fails.

The maximum number of command retry attempts does not limit the number of times that the client node can contact the server to obtain a schedule. The client node never gives up when trying to query the server for the next schedule.

Be sure not to specify so many retry attempts that the total retry time is longer than the average startup window.

If you want to have all client schedulers retry a failed attempt to process a scheduled command only twice, enter:

```
set maxcmdretries 2
```

## Setting the Amount of Time between Retry Attempts

You can specify the number of minutes the scheduler on a client node waits between retry attempts after a failed attempt to contact the server or after a scheduled command fails to process. You can use this number in conjunction with the number of command retry attempts to control when a client node contacts the server to process a failed command.

Try setting this period to half of the estimated time it takes to process an average schedule.

If you want to have the client scheduler retry failed attempts to contact the server or to process scheduled commands every 15 minutes, enter:

```
set retryperiod 15
```

---

## Managing Schedules and Associations

Task	Required Privilege Class
Define, update, copy, or delete administrative schedules	System
Define, update, copy, or delete any client schedules	System or unrestricted policy
Define, update, copy, or delete client schedules for specific policy domains	System, unrestricted policy, or restricted policy for those domains
Associate client nodes with any client schedules	System, unrestricted policy, or restricted policy
Associate client nodes with client schedules for specific policy domains	System, unrestricted policy, or restricted policy for those domains
Display information about scheduled operations	Any administrator

You can schedule both administrative commands and client operations.

Any administrator can request information about scheduled operations.

To set up an administrative command schedule on the server:

- Define the schedule
- Specify that it is an administrative type command schedule
- Specify if the schedule is to be activated or not

To set up a client schedule on the server:

- Define a schedule

**Note:** By default, schedules are client schedules. You do not have to specify the type of schedule you are defining unless it is an administrative command schedule.

- Associate client nodes with the schedule



## Defining or Updating Schedules

You can define or update schedules for both administrative commands and client operations. There are parameters on the DEFINE and UPDATE commands that apply to both administrative command and client schedules while others apply only to one type of schedule or the other. The following sections describe these parameters.

### Specifying Schedule Parameters

The following parameters on the DEFINE and UPDATE commands apply both to administrative command and client schedules:

#### Schedule name

All schedules must have a unique name, which can be up to 30 characters.

#### Initial start date, time, and day

You can specify a past date, the current date, or a future date for the initial start date for a schedule.

You can specify a start time, such as 6:00 p.m.

You can also specify the day of the week on which the startup window begins. If the start date and start time fall on a day that does not correspond to your value for the day of the week, the start date and time are shifted forward in 24-hour increments until the day of the week is satisfied.

If you select a value for the day of the week other than ANY, then depending on the values for PERIOD and PERUNITS, schedules may not be processed when you might expect. Use the QUERY EVENT command to project when schedules will be processed to ensure that you achieve the desired result.

#### Duration of a startup window

You can specify the duration of a startup window, such as 12 hours. The server must start the scheduled service within the specified duration but does not necessarily complete it within that period of time. If the schedule needs to be retried for any reason, the retry attempt must begin before the startup window elapses or the operation does not restart.

Make the window duration long enough so that all client nodes scheduled for that window have a chance to start the operation. You may have to set the window to a longer period if the number of client nodes processing the schedule is greater than the number of available scheduled sessions.

#### How often to run the scheduled service

You can set the schedule frequency based on a period of hours, days, weeks, months, or years. To have weekly backups, for example, set the period to 1 week.

#### Expiration date

You can specify an expiration date for a schedule if the services it initiates are required for only a specific period of time. If you set an expiration date, the schedule is not used after that date, but it still exists. You must delete the schedule to remove it from the database.

#### Priority

You can assign a priority to schedules. For example, if you define two schedules for one client node, and they have the same startup window, the server runs the schedule with the highest priority first. A schedule with a priority of 1 is started before a schedule with a priority of 3.

## Specifying Administrative Command Schedule Parameters

The following parameters on the DEFINE and UPDATE commands apply only to administrative command schedules:

### Administrative schedule

If you are defining an administrative command schedule, you **must** specify the command to be processed on the CMD= parameter of the DEFINE command. If you are updating an administrative command schedule, you **must** specify TYPE=ADMINISTRATIVE on the UPDATE command.

### Command

When you define an administrative command schedule, you **must** specify the complete command that is processed with the schedule. These commands are used to tune server operations or to start functions that require significant server or system resources. The functions include:

- Migration
- Reclamation
- Export and import
- Database backup

### Whether or not the schedule is active

Administrative command schedules can be active or inactive when they are defined or updated. Active schedules are processed when the specified command window occurs. Inactive schedules are not processed until they are made active by an UPDATE SCHEDULE command with the ACTIVE= parameter set to YES.

For example, to schedule the backup of the ARCHIVEPOOL primary storage pool, enter:

```
define schedule backup_archivepool type=administrative
cmd='backup stgpool archivepool recoverypool'
active=yes starttime=20:00 period=2
```

This command specifies that, starting today, the ARCHIVEPOOL primary storage pool is to be backed up to the RECOVERYPOOL copy storage pool every two days at 8:00 p.m.

To update the BACKUP\_ARCHIVEPOOL schedule, you could enter:

```
update schedule backup_archivepool type=administrative
cmd='backup stgpool archivepool recoverypool'
active=yes starttime=22:00 period=3
```

Starting with today, the BACKUP\_ARCHIVEPOOL schedule begins the backup every three days at 10:00 p.m.

## Specifying Client Schedule Parameters

The following parameters on the DEFINE and UPDATE commands apply only to client schedules:

### Domain name

A client schedule belongs to a policy domain.

### Which files or commands to process

For incremental backup operations, you can specify which file spaces to process or allow the server to perform the service based on the default client domain specified in the client user options file. Users can specify a default client domain by using the DOMAIN option in the client user options file. For information about specifying the DOMAIN option, refer to *ADSM Using the Backup-Archive Client* for the appropriate client.

For selective backup, archive, restore, and retrieve operations, you must specify the files to process.

You can use wildcard characters to select multiple files. The file spaces and file names must follow the naming conventions of the client node. Therefore, you may need to define different schedules for different platforms.

If you are scheduling a command or a macro, you must specify the entire command or the macro file name.

### Type of action

The following actions are possible:

- Perform an incremental backup
- Perform a selective backup
- Archive selected files
- Restore selected files
- Retrieve selected files
- Issue a client command
- Issue a macro

### Client options

You can specify options that are supplied to the DSMC command when the schedule is processed. You can specify most options from the client's option file. For more information, refer to the appropriate client manual.

When applicable, these options override the options specified by a client node after it has successfully contacted the server.

Do not include the following options because they have no effect on the execution of the scheduled command:

- MAXCMDRETRIES
- QUERYSCHEDPERIOD
- RETRYPERIOD
- SCHEDLOGNAME
- SCHEDMODE
- TCPCLIENTADDRESS
- TCPCLIENTPORT

To help you decide which client options and which file names or file spaces to specify when defining or updating a schedule, you can use them out during an unscheduled operation from the client node. For information about client options, refer to *ADSM Using the Backup-Archive Client* for the appropriate client.

You can define a new schedule or update an existing schedule for backing up or archiving client nodes in a specified policy domain. When you define a schedule, you assign it to a specific policy domain. You can define more than one schedule for each policy domain.

To define a schedule of incremental backups for the ENGPOLDOM policy domain, enter:

```
define schedule engpoldom engweekly action=incremental period=1 perunits=weeks
```

This command sets the incremental backup period for schedule ENGWEEKLY to 1 week to match the backup copy group frequency of the management class in the STANDARD policy set of the ENGPOLDOM policy domain.

To update the ENGWEEKLY client schedule, enter:

```
update schedule engpoldom engweekly period=5 perunits=days
```

The ENGWEEKLY schedule is updated so that the incremental backup period is now every 5 days.

## Copying Schedules

You can create a new schedule by copying an existing client or administrative schedule. When you copy a schedule, ADSM copies the following information:

- A description of the schedule
- All parameter values from the original schedule

You can then update the new schedule to meet your needs. You can copy a client schedule to another policy domain or to a newly named schedule in the same policy domain.

When you copy a client schedule, none of the client node associations are copied to the new schedule. You must associate the new schedule with client nodes before it can be used. The associations for the old schedule are not changed. See “Associating Client Nodes with Schedules” on page 153 for more information.

To copy the WINTER client schedule that belongs to policy domain DOMAIN1 to DOMAIN2 and name the new schedule WINTERCOPY, enter:

```
copy schedule domain1 winter domain2 wintercopy
```

To copy the BACKUP\_ARCHIVEPOOL administrative schedule and name the new schedule BCKSCHED, enter:

```
copy schedule backup_archivepool bcksched type=administrative
```

## Querying Schedules

When you request information about schedules, the server displays the following information:

- Schedule name
- Policy domain name
- Type of operation to be performed
- Start date and time for the initial startup window
- Duration of the startup window
- Time period between startup windows
- Day of the week on which scheduled operations can begin

The following figure shows an example of a report that is displayed after you enter:

```
query schedule engpoldom
```

Domain	* Schedule Name	Action	Start Date/Time	Duration	Period	Day
ENGPOLDOM	MONTHLY_BACKUP	Inc Bk	04/21/1994 12:45:14	2 H	2 Mo	Sat
ENGPOLDOM	WEEKLY_BACKUP	Inc Bk	04/21/1994 12:46:21	4 H	1 W	Sat

Figure 31. Example of a Schedule

## Deleting Schedules

When you delete a schedule, all associations with client nodes are also deleted. See “Associating Client Nodes with Schedules.”

To delete all schedules in the ENGPOLDOM policy domain, enter:

```
delete schedule engpoldom *
```

## Associating Client Nodes with Schedules

Client nodes process operations according to the schedule associated with the node. Client nodes can be associated with more than one schedule. However, the nodes must be assigned to the policy domain to which a schedule belongs.

After a client schedule has been defined, you can associate client nodes with it by identifying the following information:

- Policy domain to which the schedule belongs
- List of client nodes to be associated with the schedule

To associate the ENGNOD client node with the ENGWEEKLY schedule, both of which belong to the ENGPOLDOM policy domain, enter:

```
define association engpoldom engweekly engnod
```

## Querying Associations

You can display information about which client nodes are associated with a specific schedule. For example, you should query an association before deleting a client schedule.

When you query the system for information about node associations, the server returns the following information:

- Name of the schedule
- Name of the policy domain to which the schedule belongs
- Names of the clients that are currently associated with the schedule

The following figure shows the report that is displayed after you enter:

```
query association engpoldom
```

```
Policy Domain Name: ENGPOLDOM  
Schedule Name: MONTHLY_BACKUP  
Associated Nodes: MAB SSTEINER  
  
Policy Domain Name: ENGPOLDOM  
Schedule Name: WEEKLY_BACKUP  
Associated Nodes: MAB SSTEINER
```

## Deleting Associations

When you delete the association of a client node to a client schedule, the client data is no longer managed according to the schedule. However, the remaining client nodes still use the schedule.

To delete the association of the ENGNOD client from the ENGWEELY schedule, enter:

```
delete association engpoldom engweekly engnod
```

This command deletes all schedules in the ENGPOLDOM policy domain.

You might want to delete all associations to a client schedule (rather than delete the schedule) if you want to keep the schedule for future use.

## Managing Scheduled Event Records

Task	Required Privilege Class
Display information about events	Any administrator
Set the retention period for event records	System
Delete event records	System or unrestricted policy

Each scheduled administrative command and each scheduled client operation is called an *event*. All scheduled events are tracked by the server.

## Querying Event Records

To help manage event records in the database, you can request information about scheduled or completed events. You can specify a time range to limit the amount of information displayed by querying an event. To minimize the processing time when querying events:

- Minimize the time range for queried events
- For client schedules, restrict policy domains, schedules, and client node names to those for which information is required

You can request general or exception reporting queries. A general query provides information about projected and actual scheduled processes. Exception reporting provides information about scheduled processes that did not complete successfully.

If you specify a future time range, the command shows which events should occur based on current schedules.

You can query events each day to see which events were missed in the previous 24 hours. You can also query events based on the number of days that event records are retained in the database. For example, if the retention period for event records is 3 days, you can query events every 3 days.

The following figure shows the report that is displayed after you enter:

```
query event standard weekly_backup
```

Scheduled Start	Actual Start	Schedule Name	Node Name	Status
04/21/1994 18:40:00	04/21/1994 19:38:09	WEEKLY_BACKUP	GOODELL	Started

Figure 32. Example of a Report from Query Event

## Removing Event Records from the Database

You can specify how long event records stay in the database before the server deletes them. You can also manually remove event records from the database.

### Setting the Event Record Retention Period

You can specify the retention period for event records in the database. After the retention period has passed, the server automatically removes the event records from the database. At installation, the retention period is set to 10 days.

To set the retention period to 15 days, enter:

```
set eventretention 15
```

Event records are automatically removed from the database after both of the following conditions are met:

- The specified retention period has passed
- The startup window for the event has elapsed

### Deleting Event Records

Because event records are deleted automatically, you do not have to manually delete them from the database. However, you may want to manually delete event records to increase available database space.

To delete all event records written prior to 11:59 PM on June 30, 1995, enter:

```
delete event 06/30/1995 23:59
```



---

## Chapter 8. Planning Your Storage Environment

This chapter describes what you need to know before you begin managing data storage and how to plan for your storage environment. The sections listed in the following table begin at the indicated pages.

Section	Page
<b>Concepts:</b>	
How workstation files are stored	157
How the server manages storage	157
<b>Tasks:</b>	
Expanding your data storage environment	158

---

### How Workstation Files are Stored

Before you begin managing data storage, you should understand how ADSM stores workstation files in data storage.

When users back up, archive, or migrate files from client nodes to the server, ADSM:

1. Checks the management class bound to the file to determine where the file will be stored in data storage. For backed up and archived files, the server uses the storage destination defined in the applicable copy group belonging to the management class. For space-managed files, the storage destination is specified in the management class itself.

See Chapter 6, "Managing Policies" on page 113 for information on defining storage destinations, assigning copy groups to management classes, and binding management classes to user files.

2. Stores information about the files in the database. The database tracks information such as the file name, file size, owner of the file, management class attributes assigned to the file, and the location of the file in data storage.

See Chapter 4, "Managing the Database and Recovery Log" on page 63 for information on managing the database.

3. Stores the data for backed up, archived, or space-managed files on physical storage media, such as disk or tape volumes, that are assigned to storage pools.

See "Storage Pool Volumes" on page 213 for information on defining volumes to storage pools.

---

### How the Server Manages Storage

Administrators can define or update device classes. ADSM supports random and sequential access devices.

## Introduction to Devices

Each type of data storage device supported by ADSM falls into one of two categories according to the manner in which ADSM accesses the device: *random access* or *sequential access*. Typically, random access is used for high-speed devices such as disk drives, while sequential access is used for lower-speed devices such as tape drives. However, this rule of thumb is not always appropriate, and it is important to understand how ADSM manages each device.

### Random Access Devices

On random access devices, ADSM stores data by using blocks of storage that may be scattered across the available space on a volume. If some of the data becomes outdated or expires, the space occupied by that data can be reused.

ADSM provides a predefined, random-access device class, named DISK. You cannot define other random access device classes. However, you can define any number of storage pools having the DISK device class. After defining such storage pools, you must assign volumes to the storage pool. These volumes must be pre-allocated and remain fixed in size. Thus, as the need for storage grows in a storage pool that uses the DISK device class, more volumes must be preallocated and assigned to the storage pool.

### Sequential Access Devices

On sequential access devices, the server begins storing data at the beginning of a volume and appends new data after existing data. If some of the data becomes outdated or expires, the space is not reused immediately. The server can reclaim it later by using the reclamation process (see “Space Reclamation for Sequential Access Storage Pools” on page 190 for details).

## Introduction to Storage Pools

A storage pool is a collection of storage volumes that are associated with one device class (and, therefore, either random or sequential access). Many of the parameters associated with a storage pool depend on the type of access. These parameters are described in more detail in Chapter 10, “Managing Storage Pools” on page 171.

## Introduction to Storage Volumes

Administrators can define and update storage pool volumes. See Chapter 11, “Managing Storage Pool Volumes” on page 213 for information about:

- Preparing volumes for data storage
- Limiting access to volumes in a storage pool
- Defining or updating volumes in a storage pool

---

## Expanding Your Data Storage Environment

Most installations have many types of storage devices, ranging from high-performance DASDs to slower performance, but less expensive, tape devices. As an administrator, you must balance your user needs for available data with the cost of managing storage.

Use this section to evaluate your current environment to determine whether you want to define additional device classes or storage pools for your installation.

## Evaluating Your Storage Environment

Evaluate your current physical storage environment by completing the following steps:

1. Identify the storage devices and media that you have available.
2. Categorize these storage devices by their availability, performance, and storage characteristics.
3. Categorize your desired storage pools according to the devices and media types that will be used.

When evaluating your physical storage environment, identify:

- Number and type of storage devices, such as DASDs or tape devices
- Type of tape cartridges that can be used with tape storage devices
- Labeling conventions for disk and tape storage volumes

After you have evaluated your storage environment, plan the setup of your storage pools. Complete the following steps to create you plan:

1. Categorize your storage pools by device class.
2. Define the storage pools that are associated with each device class.

For example, assume that your evaluation of the physical storage environment identifies the following storage devices and media:

- 3380 DASDs
- 3390 DASDs
- 3490 Magnetic Tape Subsystem models, including two 3490 Enhanced Capacity models (3490E)
- 3480 Magnetic Tape Subsystems, including one with the improved data recording capacity (IDRC) feature
- Cartridge System Tape volumes
- Enhanced Capacity Cartridge System Tape volumes

You can categorize your storage devices by device class as shown in Table 14:

*Table 14. Categorizing Storage Devices by Device Class*

Device Class	Description
Disk	Contains storage volumes that reside on 3390 DASD
Disk	Contains storage volumes that reside on 3380 DASD
Cartridge	Contains Enhanced Capacity Cartridge System Tape volumes used with 3490E tape devices
Cartridge	Contains Cartridge System Tape volumes used with 3480 or 3490 Base tape devices

See Chapter 9, “Managing Storage Devices” on page 161 for information on defining tape device classes to support your physical storage environment.

## Considering User Needs for Recovering Data

After you have categorized your storage devices, identify availability, space, and performance requirements for user data stored on disk or tape storage. You can then assign each storage pool as a storage destination for backed up, archived, or space-managed files.

For example, an administrator determines that users in the business department have three requirements:

- Immediate access to space-managed files and to some backed up files, such as customer files and payroll accounts
- Periodic access to some archived files, such as monthly sales and inventory reports
- Occasional access to backed up or archived files that are rarely modified, such as yearly revenue reports

The administrator matches user requirements with storage devices and defines the following storage pools:

Storage Pool Name	Device Class	Volumes	Storage Destination
BACKUPPOOL	Disk	Storage volumes that reside on 3390 DASD	For a backup copy group for files requiring immediate access
SPACEMGPOOL	Disk	Storage volumes that reside on 3390 DASD	For a management class for space-managed files and requiring immediate access
ARCHIVEPOOL	Disk	Storage volumes that reside on 3380 DASD	For an archive copy group for files requiring quick, reliable access
BACKTAPE	Cartridge	Enhanced Capacity Cartridge System Tape volumes used with 3490E tape devices	For backed up data not requiring immediate access
ARCTAPE	Cartridge	Cartridge System Tape volumes used with 3480 or 3490 tape devices	For archived data not requiring immediate access

---

## Chapter 9. Managing Storage Devices

A device class represents a device type that can be used to store backup, archive, or space-managed data, or copies of storage pool data. Each storage pool is associated with just one device class. When a new device class is defined, its *device type* must be specified to indicate the device type, and thus the type of storage volumes that are to be used to store data.

For random access storage, ADSM supports only the DISK device class. However, you can define many storage pools that are categorized by the DISK device class. Although you can mix disk devices in a DISK storage pool, you should consider performance requirements for restoring or retrieving files. For example, if users need immediate access to certain files, you could define a DISK storage pool, which consists of only 3390 DASD to ensure that these files reside on faster speed devices.

ADSM supports the following device types for sequential access storage:

**CARTRIDGE** Cartridge tape drives, such as IBM 3480, 3490 and 3490E drives

**FILE** Storage volumes are files in the file system of the server machine

The information listed in the following table is presented beginning at the indicated pages.

Section	Page
<b>Tasks:</b>	
Defining and updating CARTRIDGE device classes	163
Defining FILE device classes	168
Requesting information about device classes	168
Deleting device classes	169

Most tasks presented in this chapter can be performed using either the graphical user interface or the command line interface. Table 7 on page 41 shows whether a task can be performed on the graphical user interface, the command line interface, or both.

For information about issuing ADSM commands, refer to *ADSM Administrator's Reference*. For help performing a task on the graphical user interface, see the procedure described in "Using Online Help" on page 15.

---

### Selecting a Device Type for a Device Class

Define additional CARTRIDGE device classes when your physical environment consists of different types of tape drives and storage media. However, do not mix devices that cannot read and write to the same media.

When defining CARTRIDGE device classes, consider:

- Your tape volume naming convention
- The amount of time the server retains an idle tape mount
- The maximum number of tape mounts that you have available

- The expiration date stored on tape
- The estimated capacity of tape volumes

In addition, you must consider MVS operating specific tape criteria. You should define CARTRIDGE device classes to support the different tape devices, tape recording technologies, and storage media available at your installation.

Each CARTRIDGE device class is uniquely identified by its device class name and represents a *device type*. ADSM provides one tape device type: CARTRIDGE. Examples of tape devices that support cartridge tape include the 3495 Tape Library Dataserver and the 3490 or 3480 Magnetic Tape Subsystems.

Do not mix non-compatible tape devices in a CARTRIDGE device class. For example, do not mix 3480 devices with 3490E devices because the server might not be able satisfy a tape mount request to restore user data as demonstrated by the following scenario:

1. A user backs up files, and the server attempts to store them in a tape storage pool assigned to a CARTRIDGE device class containing a 3480 and 3490E tape device.
2. Because the 3480 is already in use, the server selects the 3490E tape device and requests a tape volume that provides the highest recording format. In this case, the server selects an Enhanced Capacity Cartridge System Tape volumes on which to store the backed up files.
3. Later, the user requests to restore a file, which is located on the Enhanced Capacity Cartridge System Tape volume.
4. Because the 3490E is already in use, the server selects the 3480 tape device. However, the 3480 tape device cannot read the Enhanced Capacity Cartridge System Tape volume. Therefore, the user must wait until the 3490E device is available before the restore operation can be completed.

The device type of CARTRIDGE is *not* the same as the predefined device class of CARTRIDGE. The CARTRIDGE device class has a device type of CARTRIDGE. You can have many different CARTRIDGE device classes with the device type of CARTRIDGE.

For example, based on the example introduced in “Considering User Needs for Recovering Data” on page 160, you could define two device classes named 36TRACK and 18TRACK with the CARTRIDGE device type as shown in Table 15.

*Table 15. Naming CARTRIDGE Device Classes*

Device Class Name	Device Type	Description
36TRACK	CARTRIDGE	Supports Cartridge System Tape volumes used with 3490E tape devices
18TRACK	CARTRIDGE	Supports Cartridge System Tape volumes used with 3480 or 3490 Base tape devices

**Note:**

The 3490 Base devices include Models A01, A02, A10, A20, B02, B04, B20, B40, D41, and D42. The 3490E devices include models C1A, C2A, C10, C11, and C22

## Defining and Updating CARTRIDGE Device Classes

Task	Required Privilege Class
Define or update device classes	System or unrestricted storage

You can define and update CARTRIDGE device classes to support the tape devices and storage media available at your installation. The following parameters specify how to manage data storage operations involving the new device class:

- MOUNTLIMIT
- MOUNTWAIT
- MOUNTRETENTION
- MAXCAPACITY
- EXPIRATION
- UNIT
- COMPRESSION
- PREFIX
- PROTECTION
- RETENTION
- ESTCAPACITY

### Mount Limit

Most installations have a limited number of tape drives available at any one time. When defining CARTRIDGE device classes, you can limit the number of simultaneous tape mounts to efficiently use your tape devices. The *mount limit* specifies the maximum number of volumes that can be mounted simultaneously for a device class. Your goal in setting a tape mount limit is to guarantee that a number of tape mounts will be available to respond to a user request for restoring, retrieving, or recalling a file.

If you do not set a tape mount limit, then the server sets the tape mount limit to two. The maximum number of tape mounts per device class supported by AD SM is 256.

To determine a tape mount limit for each device class, consider:

- The number of available tape mounts for each tape device at your installation.

For example, the number of tape mounts available for some different tape devices at your installation are:

Tape Device	Number of Tape Mounts
3420 9-Track Tape Device	1
3480 Tape Subsystem	2
3490 Tape Subsystem	4

- The number and type of tape devices that you group together in a device class.

For example, if you group two 3490 Magnetic Tape Subsystem devices into one device class, then you have a total of eight available tape mounts. By limiting the tape mounts to five, you ensure that at least three mount points are available to other applications.

- Whether you want reclamation to occur.

If the mount limit is set to one, then ADSM will not be able to reclaim available space on tape volumes because it requires two tape drives in order to move readable files to another tape volume.

- How frequently migration or reclamation occurs for the tape storage pool.

For example, if the server is using all tape mounts to complete server processes, then users may wait until a tape mount becomes available before they can recover data from the tape storage pool.

See “When Files Are Migrated” on page 181 and “Space Reclamation for Sequential Access Storage Pools” on page 190 for more information.

### **Mount Wait Period**

You can use the *MOUNTWAIT* parameter to specify the maximum amount of time, in minutes, that the server waits for a manual (or operator controlled) volume mount request to be satisfied before canceling the request. The default mount wait period is 60 minutes; the maximum value for this parameter is 9999 minutes.

### **Mount Retention Period**

In addition to setting a tape mount limit, you can specify a mount retention period for each device class. The *mount retention period* is the maximum amount of time, in minutes, that the server retains a mounted tape that is not being used before it dismounts the tape. The system sets a default of 60 minutes. Mount retention time does not begin until the tape is considered idle. The following sequence occurs:

1. When the task being performed by the last client using the mounted tape has completed, that client's idletimeout clock begins.
2. When the idletimeout expires, that client session ends.
3. The mount retention time starts.
4. When the mount retention time expires, the tape is dismounted.

The duration of mount retention is the *IDLETIMEOUT* option of the last client using the tape plus the mount retention value. The *IDLETIMEOUT* option is in the server options file. The mount retention period is a parameter of the define device class command. For example, if a mounted tape volume remains idle for 60 minutes, then the server dismounts the tape volume from the tape device.

When a tape volume is needed frequently, you can improve performance by setting a longer mount retention period to avoid unnecessary tape mounts.

For manually mounted tape libraries, consider how many tape devices you group in each device class in relation to the number of tape operators who monitor those tape devices. For example, if only one tape operator supports your entire tape library on a weekend, then define a long mount retention period so that the operator is not asked to mount tapes every few minutes.

### **Maximum Capacity**

You can specify the maximum amount of data that ADSM stores in the device class. This value determines the maximum capacity of the volume. If you do not specify a value, ADSM uses a default method based on IBM 3480 and 3490E device characteristics by using IBM tapes. If the default method does not meet your needs, set the value to less than the maximum capacity of the tape to avoid



an end-of-tape situation. If ADSM detects an end-of-tape situation for the specified volume, a warning message is issued and the tape is treated as full.

### **File Expiration Date**

For each CARTRIDGE device class, you can specify an *expiration date*, which is placed in the tape label. Use the Julian date format of *yyddd* or *yyyyddd*.

ADSM passes this information to the MVS operating system, which ensures that tape volumes with existing data on them are not overwritten. This process is useful for protecting workstation files from being overwritten by another tape management system. Set this date to match the expiration date defined by the EXPDT parameter in your current JCL.

To not set an expiration date is the system default. If you specify 99365, an operator must verify whether or not the tape volume can be overwritten. If you define an expiration date, then you cannot define a retention period as described in “File Retention Period” on page 166.

### **MVS Device Unit Name**

Specify a *unit name* to define a group of tape devices that support cartridge tape, as previously defined in your JCL.

The unit name can be:

- A generic device type, such as 3480 or 3490. ADSM uses 3480 as the default.
- An esoteric name to identify a group of devices, such as CARTRIDGE or ANY-DRIVE. Use an esoteric name only when tape devices in the device class are compatible.
- A physical device. Specifying a device number is not recommended.

For example, the esoteric unit name for a group of 3480 Magnetic Tape Subsystem devices with the Improved Data Recording Capability (IDRC) feature might be *3480XF*. Or you could name a group of 3490 Magnetic Tape Subsystem devices that use Enhanced Capacity Cartridge System Tape (ECCST) as *3490LONG*.

### **Data Compression**

After you define a unit name, identify whether the device class supports the *compression* of data on tape volumes. Data compression is dependent on whether or not the tape device should exploit a hardware compaction feature, such as IDRC.

As a default, the system sets the compression attribute to yes. If your tape devices have IDRC, then files are compressed when written to cartridge tape. If your tape devices do not have IDRC, then the compression attribute is ignored and files are not compressed when written to cartridge tape.

If you plan to transport tapes to other physical sites, ensure that the second location has the capability of reading compressed data from cartridge tapes. If the alternate site cannot read compressed data, set compression to *no*.

## Tape Volume Prefix

For each CARTRIDGE device class, define a tape volume prefix. The *tape volume prefix* is the high-level qualifier of the data set name in the standard tape label. ADSM provides a default tape volume prefix named *ADSM*.

When an empty cartridge system tape is mounted, the system creates a tape label with the data set name *ADSM.BFS* to indicate that this tape volume is used by ADSM for data storage.

If you have already established a tape label data set naming convention that supports your current tape management system, consider using a tape volume prefix that conforms to your naming convention.

Determine whether you are cataloging your current tape volumes. Because the ADSM server does not catalog tape volumes, you should use the tape volume prefix to establish external security protection for cataloged tape volumes by using a security manager such as RACF.

## Protection Status

For each CARTRIDGE device class, you can set a protection attribute to *yes* or *no*, as previously defined in your JCL. The *protection* attribute updates the RACF profile to identify which users have access to cartridge volumes associated with the device class.

As a default, the system sets the protection attribute to *no*. When set to *no*, any cataloged tape volumes use the generic RACF profile. If you set the protection attribute to *yes*, RACF uses the entire data set name. This allows you to define specific protection criteria for cartridge tape volumes associated with the device class.

## File Retention Period

For cartridge tapes, you have the option of setting a retention period to specify how long files should be retained on tape. The retention period can only be used if you do not specify an expiration date.

While ADSM does not use the retention period, the server does pass the information to the MVS operating system to ensure that other tape management systems do not overwrite tape volumes that contain retained data.

For example, if you specify a retention period of 365 days, then existing tape management systems do not override files on these tapes until 365 days have passed.

Use this parameter to match the retention period specified by the RETPD parameter in your current JCL. If you do not specify a value for the retention period, ADSM sets the default to none. The maximum number of retention days that you can specify is 9999 days.

## Estimated Capacity

ADSM estimates the capacity of tape volumes assigned to storage pools categorized by device class. You can use these system defaults to plan for and monitor the use of your storage media.

The estimated capacity is used by the server to generate storage pool and storage volume reports. When volumes are written and filled, the server determines the actual capacity of the tape volume and uses this instead of the estimated capacity.

You can adjust the estimated capacity to more accurately reflect the capacity of your storage media.

Table 16 displays the estimated capacity used by ADSM for cartridge tapes.

*Table 16. Estimated Capacity for Cartridge System Tape (CST). An asterisk (\*) indicates tape formats that support compression. Because ADSM cannot determine the extent to which compression increases the capacity of a particular tape format, ADSM does not increase the estimated capacity for tape formats that support compression.*

<b>Tape Format</b>	<b>Media Type</b>	<b>Estimated Capacity</b>
3480	Not applicable	180MB
3480XF *	Not applicable	180MB
3490 Base	CST	180MB
3490B	CST	360MB
3490C *	CST	360MB
3490B	ECCST	720MB
3490C	ECCST	720MB

**Note:**

The 3490 Base devices include Models A01, A02, A10, A20, B02, B04, B20, B40, D41, and D42.

The 3490E devices include Models C1A, C2A, C10, C11, and C22.

### **Example 1: 3480 Devices**

In this example, define a device class named 3480XF to represent 3480 Magnetic Tape Subsystem devices with the IDRC feature.

Based on user requirements, set the expiration date to 99365 so that the files on these tape volumes are not expired by other tape management systems unless a tape operator overrides the expiration date. Set protection to *yes* so that MVS uses the tape volume prefix to determine which RACF profile to use with cataloging the tape volumes.

Next, use the esoteric unit name 3480 to identify that this device class supports 3480 device classes with the IDRC feature, and increase the estimated capacity to 200 MB.

Finally, use the system defaults to limit the number of tape mounts to 2, set a mount retention period of 60 minutes, set compression to *yes*, and use ADSM as the tape volume prefix.

To define the device class named 3480XF, enter:

```
define devclass 3480xf devtype=cartridge -
expiration=99365 protection=yes -
unit=3480 estcapacity=200m
```

---

## Defining FILE Device Classes

The FILE device type is used for special device classes whose storage volumes are not physical units, such as tape cartridges, but *simulated* storage volumes. Data is written sequentially into standard files in the file system of the server machine.

**Note:** Because each volume is actually a file, a volume name is a fully qualified file name string.

Additional parameters are used to instruct ADSM how to manage data storage operations involving the new device class:

### Device Type

You define the FILE device class by specifying the DEVTYPE=FILE parameter.

### Mount Limit

The mount limit value for FILE device classes is used to restrict the number of volumes (that is, files) that can be concurrently opened for access by data storage and retrieval operations. Any attempts to access more volumes than indicated by the mount limit causes the requester to wait.

See “Mount Limit” on page 163 for a discussion of the considerations necessary in determining an appropriate mount limit value for the new device class.

### Maximum Capacity

You can specify a maximum capacity value that restricts the size of volumes (that is, files) associated with a FILE device class. If you are explicitly allocating FILE volume data sets, be sure the allocated capacity matches the maximum capacity. File volumes are allocated in 16K blocks. Thus, if the maximum capacity is 50M, you should allocate 3200 blocks. A combination of primary and secondary extents may be specified to achieve the allocation. When the server detects that a volume has reached a size equal to the maximum capacity, it treats the volume as full and tries to store any new data on a different volume.

The default MAXCAPACITY value for a FILE device class is 1MB.

### Prefix

This parameter is the MVS high-level qualifier used to allocate scratch volume data sets.

---

## Requesting Information about a Device Classes

Any administrator can query for information about a device class.

To query the server to view the predefined device classes provided at installation, enter:

```
query devclass
```

Figure 33 on page 169 shows an example of a standard report for device classes.

Device Class	Device Access Strategy	Storage Pool Count	Device Type	Est. Cap. (MB)	Mount Limit	Mount Ret'n (min)	Mount Wait (min)	Format/Density
CARTRIDGE	Sequential	1	Cart.	180.0	2	60	60	Drive
DISK	Random	4						

Figure 33. Information about Device Classes

To query a server to view a detailed report for the device class named 3480XF, enter:

```
query devclass 3480xf format=detailed
```

Figure 34 shows an example of the output you receive when you request a detailed device class report from the server.

```

Device Class: 3480XF
Device Access Strategy: Sequential
Storage Pool Count: 0
Last Update by (administrator): REES
Last Update Date/Time: 03/11/1993 18:13:56
Device Type: Cartridge
Maximum Capacity (MB):
Estimated Capacity (MB): 200.0
Dataset Name Prefix: ADSM
Mount Limit: 2
Mount Retention (min): 60
Label Type: IBMSL
Expiration Date: 99365
Unit Name: 3480
Compression: Yes
Protection: Yes
Retention:

more...
```

Figure 34. A Detailed Query Device Class Report

## Deleting Device Classes

Task	Required Privilege Class
Delete device classes	System or unrestricted storage

You can delete a device class when:

- No storage pools are assigned to the device class. For information on deleting storage pools, see “Deleting a Storage Pool” on page 210.
- The device class is not being used by an export or import process.

You cannot delete the DISK device class from the server.



---

## Chapter 10. Managing Storage Pools

A storage pool is a collection of storage volumes belonging to the same device class. The storage volumes contain backed up, archived, or space-managed files.

The sections listed in the following table begin at the indicated pages.

Section	Page
<b>Concepts:</b>	
Storage pools	172
Assigning volumes in storage pools	178
Storage pool hierarchy	178
Server migration of files	181
Damaged files	185
Cache on disk storage pools	185
Collocation on sequential access storage pools	186
Space reclamation on sequential access storage pools	190
<b>Tasks:</b>	
Estimating space needs for storage pools	194
Defining or updating storage pools	194
Using copy storage pools to improve data availability	199
Backing up storage pools	200
Monitoring the use of storage pools	202
Deleting storage pools	210
Restoring storage pools	211

Most tasks presented in this chapter can be performed by using either the graphical user interface or the command-line interface. Table 8 on page 42 shows whether a task can be performed on the graphical user interface, the command line interface, or both.

For information about issuing ADSM commands, refer to *ADSM Administrator's Reference*. For help performing a task on the graphical user interface, see the procedure described in "Using Online Help" on page 15.

---

## Storage Pools

The types of storage pools are:

### Primary storage pool

Data that is backed up, archived, or migrated from a client node is stored in a primary storage pool. See “Primary Storage Pools” for details.

### Copy storage pool

When an administrator backs up a primary storage pool, the data is stored in a copy storage pool. See “Copy Storage Pools” on page 173 and “Backing Up Storage Pools” on page 200 for details.

Storage pools are categorized by the access method of the volumes contained in the storage pools:

### Random access

Associated with the DISK device class

### Sequential access

Associated with sequential device classes, for example, CARTRIDGE

ADSM has three predefined random access storage pools:

**ARCHIVEPOOL** Contains files archived from client nodes

**BACKUPPOOL** Contains files backed up from client nodes

**SPACEMGPOOL** Contains files migrated from client nodes

Although ADSM does not require a separate storage pool for migration, a separate storage pool is recommended. See “Server Migration of Files” on page 181 for more information about migration.

Figure 35 shows an example of data storage. In this example, the data storage defined for the server includes two storage pools: one disk storage pool and one tape cartridge storage pool.



Figure 35. Example of Data Storage

## Primary Storage Pools

When a client node backs up, archives, or migrates data, the data is stored in a primary storage pool. When a user tries to restore, retrieve, recall, or export file data, the requested file is obtained from a primary storage pool if possible. The primary storage pool is the default storage pool type. Primary storage pool volumes never have an access mode value of offsite. This means that the primary storage pool volumes are always located onsite.

A primary storage pool can use random access or sequential access strategy.



## Copy Storage Pools

Copy storage pools store copies of data backed up from primary storage pools. Multiple primary storage pools can use one copy storage pool. Primary storage pools can also use multiple copy storage pools if multiple copies are deemed necessary. However, it is recommended that the entire primary storage pool hierarchy be backed up to the same copy storage pool for ease of storage volume management.

**Note:** A primary storage pool file must always exist for a copy storage pool file to exist. When a primary storage pool file is deleted so are the copies of the files in the copy storage pools.

Copy storage pool volumes can have an access mode value of OFFSITE. The offsite designation allows those volumes to be moved to an offsite location and still be tracked by the database. If the ACCESS=OFFSITE parameter is not specified, the copy storage pool volumes will remain at the onsite location. Locating copy storage pool volumes offsite provides a means of recovering from an onsite disaster.

The copy storage pool concept provides a means of recovering from a disaster, data-integrity errors, or media failure. See “Restoring Storage Pools” on page 211, “Using Copy Storage Pools to Improve Data Availability” on page 199, or “Recovery from Media Loss” on page 284 for details.

Copy storage pools cannot be used as:

- The next storage pool in a migration hierarchy
- A management class destination for space-managed files
- A copy group destination for backed up or archived files

Copy storage pools do not have migration attributes. No hierarchy exists for the migration of file copies that are stored in copy storage pools.

Files in a copy storage pool do not move when primary files are moved. For example, if a file moves within a primary storage pool as a result of reclamation or is migrated to another primary storage pool, the copy file in the copy storage pool does not move. Therefore, you should use the same copy storage pool for each primary storage pool in a given migration hierarchy. This avoids the accumulation of files in different copy storage pools as the primary files are migrated. If you have multiple, non-merging, migration hierarchies, different copy storage pools can be used for primary storage pools in different hierarchies.

Normally a file is obtained from the primary storage pool whenever an attempt is made to access a file. However, ADSM attempts to access the file from a copy storage pool if the primary copy of the file cannot be obtained for one of the following reasons:

- The primary file copy has been previously marked damaged
- The primary file is stored on a volume that is UNAVAILABLE or DESTROYED
- The file is stored on an offline volume

ADSM also accesses files from a copy storage pool for certain operations (restore, retrieve, or recall of files to end users and export of file data) if the primary file is located in a storage pool that is UNAVAILABLE.

Although copy storage pools can contain only sequential access volumes and are not associated with the DISK device class, you can use disk devices for copy storage pools by specifying a device class that has a device type of FILE.

Copies in a copy storage pool are produced by using the BACKUP STGPOOL command. The copies are made incrementally while the server is operational and available to clients. Central scheduling allows for producing copies at convenient times for your installation.

## Restore Processing

ADSM provides two commands that allow an administrator to recreate files in a primary storage pool using copies in a copy storage pool:

### RESTORE STGPOOL

Command restores all files in a storage pool that have been previously identified as having data-integrity errors. These files are also known as *damaged* files. This command also restores all files on any volumes that have been designated as DESTROYED using the UPDATE VOLUME command. See “Restoring Storage Pools” on page 211 for more detailed information.

### RESTORE VOLUME

Command recreates files that reside on a volume or volumes in the same primary storage pool. This command can be used to recreate files for one or more volumes that have been lost or damaged. See “Restoring Storage Pool Volumes” on page 230 for more detailed information.

ADSM uses database information to determine which files should be restored for a volume or storage pool, so restore processing does not require that the original volumes be accessed. For example, if a primary storage pool volume becomes damaged, the RESTORE VOLUME command could be used to recreate files that were stored on that volume, even though the volume itself is not readable. However, if the administrator were to delete the damaged files with DISCARDDATA=YES, the database reference to the files on the primary storage pool volume and all references to copies of the files on copy storage pool volumes, would be removed from the database. It would not be possible to restore those files.

Restore processing obtains files from a copy storage pool and stores these files on new primary storage pool volumes. Database references to files on the original primary storage pool volumes are then deleted. If a primary storage pool volume becomes empty because all files that were stored on that volume have been restored, the empty volume is automatically deleted from the database.

To facilitate restore processing of entire volumes, the DESTROYED volume access mode is used. If a volume has an access mode of DESTROYED, that volume will not be mounted for either read or write access. This mode is used to designate primary volumes for which files are to be restored. The RESTORE VOLUME command automatically changes the access mode of specified volumes to DESTROYED using a volume list provided by the command syntax. The RESTORE STGPOOL command requires the administrator to update volumes to DESTROYED before the RESTORE STGPOOL command is issued.

The DESTROYED designation for volumes is important during restore processing, particularly when the RESTORE STGPOOL command is used to restore a large number of primary storage pool volumes after a major disaster:

- It provides a means of designating those volumes that need to be restored. If some volumes are known to be usable after a disaster, the access state of the usable volumes should not be set to DESTROYED so they will not be restored.
- Once the administrator has identified the primary volumes to be restored, and has changed the access mode of these volumes to DESTROYED, new volumes can be added to the storage pool. The new volumes are used to contain the files as they are restored from the copy storage pool volumes, and can also be used for storage of new files that may be backed up, archived, or migrated by the end users.
- The designation of DESTROYED volumes allows ADSM to keep track of the files that still need to be restored from copy storage pools. If restore processing should be prematurely terminated for any reason, processing could be resumed and only the files that still reside on DESTROYED volumes would need to be restored.

### Expiration Processing

When filespaces are deleted, backup files are versioned off, or archive files pass their archive retention period, these files are expired from the ADSM database. When files expire, any copies of those files made in copy storage pools are also removed from the database. If backup policies are setup appropriately, the need to recover an expired file should be a rare occurrence.

If this need occurs, expired files can be recovered by:

1. Restoring the database to a point in time prior to file expiration.
2. Using a primary or copy storage pool volume that has not been rewritten and contains the expired file data at the time of database backup.

To ensure expired files can be recovered, the REUSEDELAY parameter for copy storage pools (on the DEFINE STGPOOL and UPDATE STGPOOL command) should be set at least as long as your oldest database backup. When backing up primary storage pools, the REUSEDELAY parameter for the primary storage pools should be set to 0, to efficiently reuse primary scratch volumes.

### Migration Processing

Migration of files between primary storage pools does not affect copy storage pool files. Copy storage pool files do not need to be moved when primary storage pool files move.

For example, suppose a copy of a file is made while it is in a disk storage pool. The file then migrates to a primary tape storage pool. If you then backup the primary tape storage pool to the same copy storage pool, a new copy of the file is not needed. ADSM knows it already has a valid copy of the file.

**Note:** The migration of files is not allowed for copy storage pools.

See “Server Migration of Files” on page 181 for more information about migration.

## Reclamation and MOVE DATA Command Processing

Reclamation and MOVE DATA command processing of primary storage pool volumes does not affect copy storage pool files.

Reclamation and MOVE DATA command processing of volumes in copy storage pools is similar to that of primary storage pools. However, there are a couple differences between primary and copy storage pool processing. One difference is that most volumes in copy storage pools may be set to an access mode of OFFSITE, making them ineligible to be mounted. During reclamation or MOVE DATA command processing, valid files on offsite volumes are copied from the original files in the primary storage pools. In this way, valid files on offsite volumes are copied without having to mount these volumes. These new copies of the files are written to another volume in the copy storage pool.

Another difference between primary and copy storage pool processing is in the MOVE DATA function. A MOVE DATA command on a primary storage volume can be directed to any primary storage pool. A MOVE DATA command on a copy storage volume can not be directed to another storage pool. It can only go to the same copy storage pool.

See “Moving Files from One Volume to Another Volume” on page 225 for more information about moving data.

Reclamation of copy storage pool volumes should be done periodically to allow reuse of partially filled volumes that are offsite. Reclamation can be done automatically by setting the reclaim threshold for the copy storage pool. One special consideration occurs when doing it automatically for copy storage pools.

Suppose you plan to make daily storage pool backups to a copy storage pool, then mark all new volumes in the copy storage pool as OFFSITE and send them to the offsite storage location. This strategy works well with one consideration if you are using automatic reclamation.

Each day’s storage pool backups will create some number of new copy storage pool volumes, the last one being only partially filled. If this partially filled volume is emptier than the reclaim percentage, this volume becomes eligible for reclamation as soon as you mark it OFFSITE. This would cause a new volume to be created with the same files on it. The volume you take offsite would be empty according to the ADSM database. If you don’t recognize what is happening, you could perpetuate this process by marking the new partially filled volume OFFSITE.

This consideration with automatic reclamation could be resolved by keeping partially filled volumes onsite until they fill up. However, this would mean a small amount of your data would be without an offsite copy for another day.

For this reason, it is recommended you control copy storage pool reclamation by using the default value of 100. This turns reclamation off for the copy storage pool. You can initiate reclamation at desired times by resetting the reclaim threshold with the UPDATE STGPOOL command. To monitor offsite volume utilization and help you decide what reclamation threshold to use, enter the following:

```
query volume * access=offsite format=detailed
```

Depending on your data expiration patterns, you may not need to do reclamation of offsite volumes each day. You may choose to perform offsite reclamation on a less frequent basis.

When you do perform offsite reclamation, it is recommended that you turn on reclamation for copy storage pools during your storage pool backup window and before marking copy storage pool volumes as OFFSITE. Next, turn off reclamation and then mark any newly created volumes as OFFSITE.

It is also recommended that you set the REUSEDELAY parameter for your copy storage pool to be at least as long as the oldest database backup you intend to keep. This will ensure that reclaimed volumes are retained long enough to guarantee the recovery of expired files. Offsite volumes that you see are in the PENDING state are empty but are awaiting release based on the REUSEDELAY value.

See “Space Reclamation for Sequential Access Storage Pools” on page 190 for more information about reclamation.

### **Collocation on Copy Storage Pools**

There are special considerations when using collocation on copy storage pools. Primary and copy storage pools perform different recovery roles. Direct client recovery is typically done from the primary pools while copy storage pools are usually used to recover the primary pool data. In a disaster where both clients and the server are lost, the copy storage pool volumes will probably be used directly to recover clients. The types of recovery scenarios that are of most concern to you will help to determine whether to use collocation on your copy storage pools.

Another consideration is that collocation on copy storage pools will result in more partially filled volumes and potentially unnecessary offsite reclamation activity.

Collocation typically results in a partially filled sequential volume for each client. This may be acceptable for primary storage pools because these partially filled volumes remain available and can be filled during the next migration process. However, for copy storage pools this may be unacceptable because the storage pool backups are usually made to be taken offsite immediately. If you use collocation for copy storage pools, you will have to decide between:

- Taking more partially filled volumes offsite thereby increasing the reclamation activity when the reclamation threshold is lowered or reached.

**or**

- Leaving these partially filled volumes onsite until they fill and risk not having a offsite copy of the data on these volumes.

With collocation disabled for a copy storage pool, typically there will be only a single partially filled volume after storage pool backups to this copy storage pool are complete.

Careful consideration should be given before using collocation for copy storage pools. Even customers using collocation for their primary storage pools may wish to disable collocation for copy storage pools. One example of when collocation on copy storage pools may be desirable is when you have few clients, but each of them has large amounts of incremental backup data each day.

See “Collocation on Sequential Access Storage Pools” on page 186 for more information about collocation.

---

## Assigning Volumes in Storage Pools

Volumes are assigned differently depending on whether they are stored in random access storage pools or sequential access storage pools.

### Assigning Random Access Storage Pool Volumes

Volumes in random access storage pools must be predefined.

### Assigning Sequential Access Storage Pool Volumes

You can predefine volumes in a sequential access storage pool or you can specify that ADSM dynamically acquire scratch volumes. You can also use a combination of predefined and scratch volumes.

Use predefined volumes when you want to control which volumes are used in the storage pool. This process may be useful when you want to establish a volume naming scheme for ADSM volumes. See Chapter 11, “Managing Storage Pool Volumes” on page 213 for information about defining volumes.

Use scratch volumes when you want to allow ADSM to dynamically acquire a volume when needed and dynamically delete the volume when it becomes empty. For example, you might want to use scratch volumes to avoid the burden of explicitly defining all of the volumes in a given storage pool.

Scratch volumes in a copy storage pool with an access value of offsite are handled differently than scratch volumes with other access values in a primary storage pool. If an offsite volume becomes empty, it is not immediately returned to the scratch pool. This prevents the volumes from being deleted from the database and makes it easier to determine which volumes should be returned to the onsite location. The volume is not returned to the scratch pool until the access value is changed to READWRITE, READONLY, or UNAVAILABLE. This allows the administrator to query ADSM for empty offsite copy storage pool volumes and return them to the onsite location.

---

## Storage Pool Hierarchy

Consider using multiple levels of primary storage pools to form a storage hierarchy. For example, assume that your fastest devices are disks, but space on these devices is scarce. You also have tape drives, which are slower to access, but have much greater capacity. You can define a hierarchy so that files are initially stored on the fast disk volumes in one storage pool, and provide clients with quick response to backup and recall requests. However, as the disk storage pool becomes full, data can be migrated, or moved, to tape volumes in a different storage pool. Migrating files to sequential storage pool volumes is particularly useful because all the files for a node are migrated together and organized in a more orderly way. This is especially helpful if collocation is not enabled.

When defining or updating a storage pool, you establish a hierarchy by identifying the storage pool to which data will be migrated, or moved, if the original storage pool is full or otherwise unavailable.

One of the most important considerations is the number of copy storage pools you need. This will depend on the hierarchies you have set up with your primary storage pools and what type of disaster recovery protection you wish to implement. For most cases, a single copy storage pool could be used for backup of all primary storage pools.

Multiple copy storage pools may be needed to handle particular situations, including:

- Special processing of certain primary storage hierarchies (for example, archive pools or storage pools dedicated to priority clients)
- Creation of multiple copies for multiple locations (for example, to keep one copy onsite and one copy offsite)
- Rotation of full storage pool backups (See “Full Storage Pool Backups” on page 201 for more information.)

It is strongly recommended that all primary storage pools that are linked to form a storage hierarchy use the same copy pool for backup. If this is done, then a file that is copied does not need to be recopied when it migrates to another primary storage pool.

**Note:** You cannot establish a chain of storage pools that leads to an endless loop. For example, you cannot define StorageB as the *next* storage pool for StorageA, and then define StorageA as the *next* storage pool for StorageB.

## Storage Considerations

A useful guideline for how much primary disk storage should be dedicated for the staging of client data is: enough storage to handle one night’s worth of the clients’ incremental backups. This is not always feasible but it is a good guideline. This guideline has even more value when considering storage pool backups.

If one day’s worth of client incrementals can be kept in a disk storage pool, backup copies of these files can be made while they are residing on disk, before they are migrated to sequential media. This saves many mount requests while performing your storage pool backups. Backing up your storage pools before migration processing will allow you to copy as many files as possible while they are still on disk. The primary sequential storage pools must still be backed up to catch any files that might have been missed while on disk (for example, large files that went directly to sequential media).

See “Estimating Space Needs for Storage Pools” on page 192 for more information about storage pool space.

## How ADSM Stores Files in a Storage Pool Hierarchy

This section explains how the server selects and accesses a primary storage pool. This information can help you estimate the amount of space required for each storage pool in the hierarchy.

When a user backs up, archives, or migrates a file from a client node to the server, the server looks at the management class that is bound to the file to determine in which storage pool to store the file. The server then checks the storage pool to determine the following:

- If it is possible to write file data to the storage pool (access mode)

- What the maximum file size allowed is in the storage pool
- What the high migration threshold is for the storage pool
- If sufficient space is available on the available volumes in the storage pool
- What the next storage pool used is, if any of the previous conditions prevent the file from being stored in the storage pool being checked

Based on these factors, the server determines if the file can be written to that storage pool or the next storage pool in the hierarchy. An example of how this might work follows:

Assume a company has a storage pool hierarchy as shown in Figure 36.




Figure 36. Storage Hierarchy, Read/Write Access, and Maximum File Size

The storage pool hierarchy consists of four storage pools:

- FASTDISKPOOL** The top of the storage hierarchy. It contains fast disk volumes for storing data.
- SLOWDISKPOOL** The next storage pool in the hierarchy. It contains slower disk volumes.
- FASTTAPEPOOL** The next storage pool in the hierarchy. It contains tape volumes accessed by high-performance tape drives.
- SLOWTAPEPOOL** The last storage pool in the hierarchy. It contains tape volumes accessed by lower-performance tape drives.

Assume a user wants to archive a 5MB file named *FileX*. *FileX* is bound to a management class that contains an archive copy group whose storage destination is SLOWDISKPOOL, see Figure 36.

When the user archives the file, the server determines where to store the file based on the following decision making process:

1. The server selects SLOWDISKPOOL because it is the specified archive storage destination.
2. Because the access mode for SLOWDISKPOOL is read/write, the server checks the maximum file size allowed in the storage pool.



3. The maximum file size allowed in SLOWDISKPOOL is 3MB. FileX is a 5MB file and therefore cannot be stored in SLOWDISKPOOL. The server searches for the next storage pool in the storage hierarchy.
4. The server checks the access mode of FASTTAPEPOOL, which is the next storage pool in the storage hierarchy.
5. The access mode for FASTTAPEPOOL is read-only. Therefore, the file cannot be stored in FASTTAPEPOOL, and the server searches for the next storage pool in the storage hierarchy.
6. The next storage pool in the storage hierarchy is SLOWTAPEPOOL, and the server checks the access mode.
7. The access mode for SLOWTAPEPOOL is read/write. The server then checks the maximum file size allowed in the storage pool.
8. Because SLOWTAPEPOOL is the last storage pool in the storage hierarchy, no maximum file size is specified. Therefore, if there is available space in SLOWTAPEPOOL, FileX can be stored in it.

---

## Server Migration of Files

ADSM provides an automatic migration facility to maintain free space in a primary storage pool. For example, stored data on a random access disk storage pool can be set to migrate to a less expensive sequential access storage pool. Threshold parameters must be defined before migration can take place.

## Migration Thresholds for Disk Storage Pools

When you define or update a storage pool, set migration thresholds to specify when the server should begin migrating, or moving, data to the next storage pool in the storage hierarchy. This process helps to ensure that there is sufficient free space in the storage pools at the top of the hierarchy.

You can use the DEFINE STGPOOL command defaults for the migration threshold, or you can change the threshold values to identify the maximum and minimum amount of space for a storage pool. See “Defining a Primary Storage Pool” on page 194 for more information about migration thresholds.

Before you define migration thresholds, you should understand how the server determines when to migrate files, and how it chooses which files to migrate. Then you can determine migration thresholds for both disk and sequential access storage pools.

For disk storage pools, migration thresholds can be set lower when cache is enabled. See “The Use of Cache on Disk Storage Pools” on page 185 for information about setting the CACHE parameter.

### When Files Are Migrated

Migration from disk storage pools is performed by node, and nodes with the largest file space are migrated first. When the high-migration threshold is reached in a disk storage pool, the server searches for the client file space that consumes the most space in the storage pool. Then the server migrates *all* files from *every* file space belonging to that client. After doing so, the server checks the low-migration threshold for the storage pool to determine if the migration process should be stopped. If, after migrating the given client’s files, the amount of used space in the

storage pool drops below the low migration threshold, migration ends. If not, another client's file spaces are chosen by using the same criteria as described above, and the migration process continues.

For example, Table 17 displays information contained in the database that is used by the server to determine which files to migrate.

<i>Table 17. Database Information on Files Stored in DISKPOOL</i>	
<b>Client Node</b>	<b>Data Storage Used by File Space</b>
TOMC	TOMC/C = 200MB TOMC/D = 100MB
HTANG	HTANG = 50MB
PEASE	PEASE/home = 150MB PEASE/temp = 175MB

In addition, Figure 37 displays the migration thresholds defined for the disk storage pool *DISKPOOL* and the tape storage pool *TAPEPOOL*.

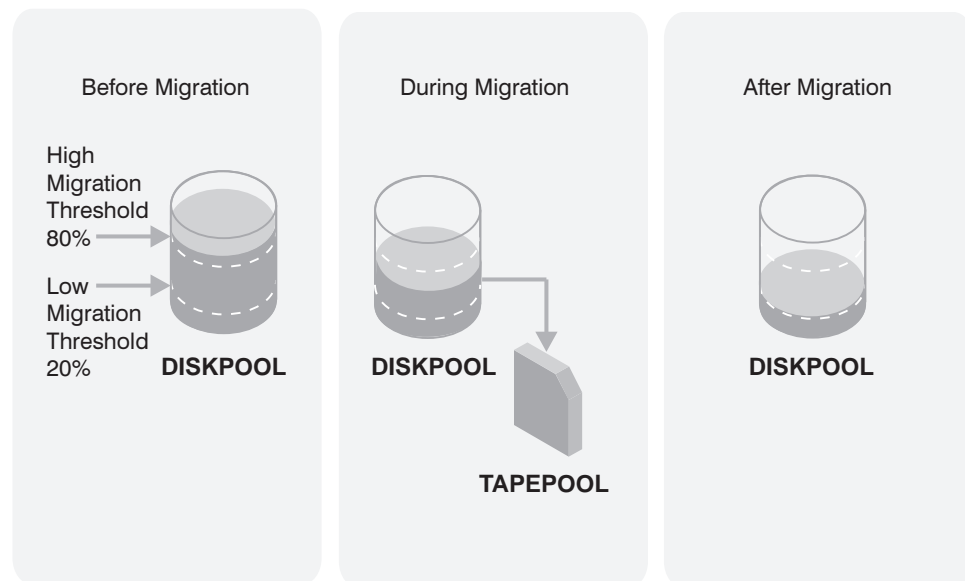


Figure 37. Setting High- and Low-Migration Thresholds

When the amount of data in DISKPOOL reaches 80%, the server performs the following tasks:

1. Determines that the TOMC/C file space is taking up the most space in the DISKPOOL storage pool.
2. Locates all data belonging to node TOMC stored in DISKPOOL. In this example, node TOMC has backed up or archived files from file spaces TOMC/C and TOMC/D stored in the DISKPOOL storage pool.
3. Migrates all data from TOMC/C and TOMC/D to the next available storage pool. In this example, the data is migrated to the tape TAPEPOOL storage pool.

The server migrates all of the data from both file spaces belonging to node TOMC, even if the occupancy of the storage pool drops below the low-migration threshold before the second file space has been migrated. For example, by moving data from file spaces TOMC/C and TOMC/D, the space utilization of DISKPOOL could drop below 20%.

If the cache option is enabled, the space utilization of DISKPOOL remains consistent because the server leaves cached copies of migrated files on the storage pool. Cached copies remain on disk storage until space is needed for new files. By using cache, you can improve the retrievability of files. When this option is used, the occupancy value which the server compares against the migration thresholds does not include space occupied by cached copies of files.

After all files that belong to TOMC are migrated to the next storage pool, the server checks the low-migration threshold. If the low-migration threshold has not been reached, then the server begins migrating files belonging to the remaining client node with the file space that is using the most space in the storage pool. In this example, the server migrates *all* files that belong to the client node named PEASE to the TAPEPOOL storage pool.

After all the files that belong to PEASE are migrated to the next storage pool, the server checks the low-migration threshold again. If the low-migration threshold has been reached or passed, then migration ends.

### **Appropriate Migration Threshold Values**

Setting migration thresholds for disk storage pools ensures sufficient free space on faster speed devices, without having migration occur so frequently that the device is unavailable for other use.

To calculate the high-migration threshold, consider:

- The amount of storage capacity provided for each storage pool
- The amount of free storage needed for users to add to existing files, without having migration occur

To calculate the low-migration threshold, consider:

- The amount of free disk storage space needed for normal daily processing
- Whether to use cache on disk storage pools to improve the retrievability of data
- How frequently you want migration to occur, based on the availability of sequential access storage devices and mount operators

### **Immediate User Access to Files on Disk Storage**

Caching is a good method of providing immediate access to files on disk storage. However, if you need to prevent files on disk storage from migrating to other storage pools. Use any of the following methods:

- Do not define the *next* storage pool
- Set the high-migration threshold to 100%

By setting the high migration threshold to 100%, you can still define the *next* storage pool in the storage hierarchy, so that large files will be moved to a storage pool that supports large files as defined by the maximum file size.

## Migration Thresholds for Sequential Access Storage Pools

Migration from sequential storage pools is performed by volume. This is done to minimize the number of mounts for source volumes. Sequential volumes selected for migration are those that were least recently referenced.

There is no straightforward way to selectively transfer data for a specific node from one sequential storage pool to another. If you know the volumes on which a particular node's data is stored, you can use the MOVE DATA command to move files from selected volumes to the new storage pool.

While you can define or update migration thresholds for sequential access storage pools, you probably will not perform this type of migration on a regular basis. This type of operation, such as tape-to-tape migration, has limited benefits compared to, disk-to-tape migration and requires at least two tape drives.

However, you may find it necessary to migrate data from one sequential access storage pool to another. For example, if you install a different tape drive or you want to move tape volumes from an automatic tape library to shelf volumes, then migration from a sequential access storage pool may be appropriate.

When defining migration criteria for sequential access storage pools, consider:

- The capacity of the volumes in the storage pool
- The time required to migrate data to the next storage pool
- The speed of the underlying devices
- The time required to mount media, such as tape volumes, into drives
- Whether operator presence is required

If you decide to migrate data from one sequential access storage pool to another, ensure that:

- Two drives (mount points) are available.
- The next storage pool in the storage hierarchy has read/write access.

For information about setting an access mode for sequential access storage pools, see "Defining a Primary Storage Pool" on page 194.

- Collocation is set similarly in both storage pools. For example, if collocation is set to *yes* in the first storage pool, then collocation should be set to *yes* in the subordinate storage pool.

For information about enabling or disabling collocation for sequential access storage pools, see "Collocation on Sequential Access Storage Pools" on page 186.

- You have sufficient staff available to handle any necessary media mount and dismount operations, because the server attempts to reclaim space from sequential access storage pool volumes before it migrates files to the next storage pool.

For information about setting a reclamation threshold for tape storage pools, see "Space Reclamation for Sequential Access Storage Pools" on page 190.

If you want to limit migration from a sequential access storage pool to another storage pool, set the high-migration threshold to a high percentage, such as 95%.

---

## Damaged Files

If a user experiences a data-integrity error during an attempt to restore, retrieve, or inspect a file during an AUDIT VOLUME operation, the file is marked as damaged. The file is marked as damaged in the database whether the file is in a primary or copy storage pool. If the same file is stored in other storage pools, the status of those file copies is not changed.

If a user attempts to restore, retrieve, or recall a file that is marked as damaged, ADSM will send the user a copy of the file if a copy is available on an onsite copy storage pool volume.

Files that are marked as damaged are not:

- Restored, retrieved, or recalled
- Moved by migration, reclamation, or the MOVE DATA command
- Backed up during a BACKUP STGPOOL operation (if the primary file is damaged)
- Restored during a RESTORE STGPOOL or RESTORE VOLUME operation (if the backup copy in a copy storage pool is also damaged)

The AUDIT VOLUME command can be used to reset the damaged status of files if the volume in which the file is stored is audited without detecting data-integrity errors. This allows the administrator to correct situations when files are marked damaged as a result of a temporary hardware problem such as a dirty tape head.

If a primary file copy is marked as damaged and a usable backup copy exists in a copy storage pool, the primary file can be recreated using the RESTORE VOLUME or RESTORE STGPOOL command.

ADSM provides commands that allow administrators to identify files that are marked as damaged:

- The RESTORE STGPOOL command displays the name of each volume in the restored storage pool that contains one or more damaged primary files. This command can be used with the preview option to identify primary volumes with damaged files without actually performing the restore operation.
- The QUERY CONTENT command provides a DAMAGED option that allows the administrator to display damaged files on a specific volume.

---

## The Use of Cache on Disk Storage Pools

When defining or updating disk storage pools, you can enable or disable cache. When cache is enabled, the migration process leaves behind duplicate copies of files on disk after the server migrates these files to subordinate storage pools in the storage hierarchy. The copies remain in the disk storage pool, but in a *cached* state, so that subsequent retrieval requests can be satisfied quickly. However, if space is needed to store new data in the disk storage pool, the space occupied by cached files can be immediately reused for the new data.

By default, the system enables caching for each disk storage pool. However, this option can be changed when a storage pool is defined or updated.

## Considerations for Using Cache

Using cache improves the retrievability of files, because a copy of the file remains on fast disk storage after the primary file is migrated.

When cache is used and migration occurs, the server migrates backed up or archived files, but leaves cached copies in the disk storage pool. The cached copies remain in the disk storage pool until space is needed for new files.

When space is needed, the server reclaims space by writing over the cached files, that have the oldest retrieval date and occupy the largest amount of disk space. For example, if File A was last retrieved on 04/16/95 and File B was last retrieved on 06/19/95, then File A is deleted to reclaim space before File B.

## Considerations for Not Using Cache

Set cache to *no* if you have limited database space because the server has to keep track of both the cached copy of the file and the new copy in the subordinate storage pool.

When cache is not used and migration occurs, the server migrates the backed up or archived files to the next storage pool and erases the file from the disk storage pool.

If you disable cache, you may want to set higher migration levels for the disk storage pool to prevent migration from occurring too frequently.

---

## Collocation on Sequential Access Storage Pools

*Collocation* is a process in which the server attempts to keep all files belonging to a client file space on a minimal number of sequential access storage volumes.

To have ADSM collocate data when files from different client nodes are mixed in the same storage pool, set collocation to *yes* when you define or update a sequential storage pool. By using collocation, you reduce the number of volume mount operations required when users restore, retrieve, or recall many files from the storage pool. Collocation thus improves access time for these operations. Figure 38 shows an example of collocation enabled, with three clients having separate volumes assigned for each client.




Figure 38. Example of Collocation Enabled

When collocation is disabled, the server attempts to use all available space on each volume before selecting a new volume. While this process provides better utilization of individual volumes, user files can become scattered across many volumes. Figure 39 on page 187 shows an example of collocation disabled, with three clients sharing space on a volume.




Figure 39. Example of Collocation Disabled

When users want to restore, retrieve, or recall a large number of files, media mount operators may be required to mount more volumes to recover user data. The system default is to not use collocation.

To determine whether or not to use collocation, consider:

- The amount of time available for backup processing  
If you have limited time for backup, disable collocation because with collocation you have more media mounts.
- The amount of time required to access a particular sequential access storage volume  
The access time depends heavily on the type of media involved in the operation. For example, if the underlying device is a tape, the access time is long, because volumes must be inserted into the appropriate type of drive via manual intervention or robotic load. However, if the device type of the device class associated with the storage pool is FILE, then the storage volumes can typically be accessed very quickly, and without manual intervention.

- Whether users need to be able to restore or retrieve a large number of files within a short period of time

When users need to restore or retrieve a large number of files, enable collocation. Without collocation, your ability to recover files for users might be delayed because:

- More than one user's files can be stored on the same sequential access storage volume.

For example, if two users attempt to recover a file that resides on the same volume, the second user will be forced to wait until the first user's files are recovered.

- A user's files can be spread across multiple volumes, requiring additional media mounts and dismounts by operators.
- How you want the server to utilize storage space
 

When collocation is enabled, the server attempts to use all available tape volumes to separate user files before it uses all available space on every tape volume.

When collocation is disabled, the server attempts to use all available space on each tape volume before it selects the next tape volume.
- Whether you have sufficient personnel to manage media mounts during backup, archive or client migration operations
 

While collocation helps to reduce the number of mount operations during recovery, operators may experience:

  - More mounts when user files are backed up, archived, or migrated from client nodes directly to sequential access volumes
  - More mounts during reclamation or migration
  - Additional handling of sequential access volumes because the volumes might not be fully used

To reduce the number of media mounts and to use space on sequential volumes more efficiently, you can:

- Define a storage pool hierarchy that requires backed up, archived, or client-migrated files to be stored initially in disk storage pools.
 

When files are migrated from a disk storage pool, the server attempts to migrate all files belonging to the client node which is using the most disk space in the storage pool. This process works well with the collocation option because the server tries to place all of the files from a given client on the same sequential access storage volume.
- Use scratch volumes for sequential access storage pools to allow the server to select new volumes for collocation.

## How the Server Selects Volumes with Collocation Enabled

When collocation is enabled and user files are backed up, archived, or migrated to sequential access storage, the server attempts to select a volume that already contains files from file spaces belonging to the client node.

If no such volume exists, the server attempts to select an empty volume. The server first selects volumes that have been explicitly defined in the storage pool. If no predefined volumes exist, but scratch volumes are supported for the storage pool, the server attempts to select a scratch volume.

If no empty volume exists and no scratch volume can be obtained, the server selects the emptiest volume that already contains data.

When the server needs to continue to store data on a second volume, it uses the following selection order to acquire additional space:

1. An empty predefined volume
2. An empty scratch volume
3. A volume on which other user files are already stored
4. A volume that has the most available free space



5. Any available volume in the storage pool

Through this selection process, the server attempts to provide the best use of individual volumes without mixing user files on multiple volumes. For example, Figure 40 shows that volume selection is *horizontal*, where all available volumes are used before all available space on each volume is used. A, B, C, and D represent data for a single node.




Figure 40. Using All Available Sequential Access Storage Volumes with Collocation Enabled

## How the Server Selects Volumes with Collocation Disabled

When collocation is disabled, the server attempts to use all available space in a storage volume before it accesses another volume. When storing client files in a sequential access storage pool where collocation has been disabled, the server first attempts to select a previously used sequential volume with available space.

If none exists, the server selects the volume that contains the most data so that each volume is fully utilized. If no partially full volume exists, the server selects an empty volume.

When the server needs to continue to store data on a second volume, it attempts to select an empty volume. If none exist, the server attempts to select any remaining available volume in the storage pool.

Figure 41 shows that volume utilization is *vertical* when collocation is disabled. In this example, fewer volumes are used because the server attempts to use all available space by mixing user files on individual volumes.

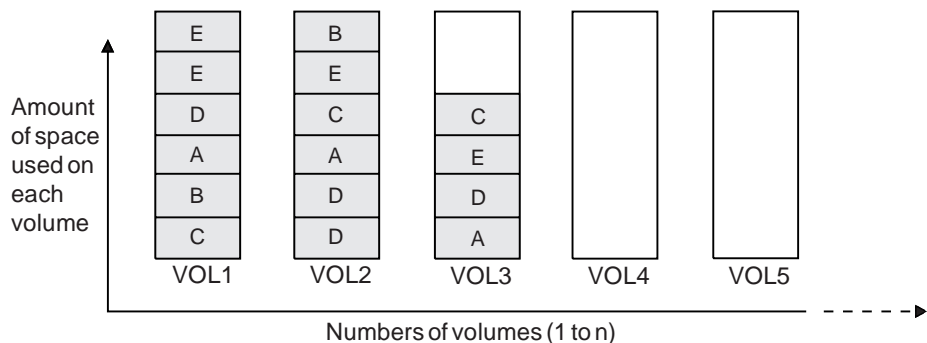


Figure 41. Using All Available Space on Sequential Volumes with Collocation Disabled

## Space Reclamation for Sequential Access Storage Pools

When files on a sequential volume become obsolete due to aging or version limits, the space on the volume becomes reclaimable. When the percentage of reclaimable space exceeds a specified level, the *reclamation threshold*, the server begins space reclamation for the volume.

During space reclamation, the server copies active files from the candidate volume to other volumes in the storage pool. For example, Figure 42 shows the active files from tapes 1, 2, and 3, being consolidated on tape 4.




Figure 42. Tape Reclamation

After all readable files have been moved to other volumes, one of the following actions is taken for the candidate volume:

- Reused, if the volume has been defined to the storage pool
- Deleted, if the volume has been acquired as a scratch volume

Volumes in a copy storage pool are reclaimed in the same manner as a primary storage pool with the exception of *offsite* volumes.

## Reclamation of Offsite Volumes

Primary storage pool volumes with an access value of `readwrite` or `readonly` must be full before reclamation can occur. Reclamation of copy storage pool volumes, with the access value of `offsite`, can occur even if the volume is not full. Volumes with the access value of `offsite`, are eligible for reclamation if the amount of empty space on a volume exceeds the reclamation threshold for the copy storage pool. A partially filled volume can be moved offsite and be reclaimed if it contains more empty space than is allowed by the reclamation threshold. When an offsite volume is reclaimed, the files on the volume are rewritten to a `readwrite` volume. Effectively these files are moved back to the onsite location, but may be obtained from the offsite volume after a disaster if a database backup is used that references the files on the offsite volume. The default reclamation threshold for copy storage pools is 100%, which means that reclamation is not performed.

Reclamation processing for offsite volumes is performed as follows:

1. The server determines which files are still active on the volume to be reclaimed.
2. These active files are obtained from a primary storage pool (or from another onsite volume of a copy storage pool, if necessary).
3. The active files are written to one or more new volumes in the copy storage pool and the database is updated.
4. A message is issued indicating that the offsite volume was reclaimed.

## Reclamation Threshold

If the reclamation threshold is low, the server tries to reclaim space occupied by obsolete files more frequently. Frequent reclamation optimizes the use of a sequential access storage pool's space.

However, each reclamation process requires *at least* two volume mounts. This process could significantly increase the manual intervention needed to mount sequential volumes.

If the reclamation threshold is 50% or lower, ADSM may not be able to combine the usable files from multiple volumes onto a single new volume. At least two drives in the same device class are required for reclamation. For more information about reclamation, see "Mount Limit" on page 163. There must be a sufficient number of volumes, drives (if appropriate), and mount operators (if appropriate) to handle frequent reclamation requests.

If the reclamation threshold is high, reclamation occurs less frequently. A high reclamation threshold is useful if manual volume mounts are required and the operations staff is at a minimum. Setting the reclamation threshold to 100% prevents reclamation from occurring at all.

## How Collocation Affects Reclamation

If collocation is enabled and reclamation occurs, the server tries to reclaim each user's files onto a minimal number of volumes. Therefore, if the volumes are manually mounted, the mount operators must:

- Observe that a tape volume may be rewound more than once if the server completes a separate pass to move each client's data.

- Mount and dismount multiple volumes to allow the server to select the most appropriate volume on which to move each client data. The server tries to select a volume in the following order:
  1. A volume that already contains files belonging to the client node
  2. An empty volume
  3. The volume with the most available space
  4. Any available volume

If collocation is disabled and reclamation occurs, the server tries to move usable data to new volumes by using the following volume selection criteria:

- The volume that contains the most data
- Any partially full volume
- An empty predefined volume
- An empty scratch volume

## Reclamation in a Single Drive Library

Reclamation of volumes in a single drive library are performed by using the MOVE DATA command. If the target storage pool is higher in the storage pool hierarchy than the original storage pool, the moved data will migrate back into original storage pool and be written to a new volume. The original storage pool volume is then reclaimed.

---

## Estimating Space Needs for Storage Pools

This section provides guidelines for estimating the initial storage space required for your installation. It assumes the use of the following the default random access (disk) storage pools provided by ADSM: following:

- BACKUPPOOL for backed up files
- ARCHIVEPOOL for archived files
- SPACEMGPOOL for files migrated from client nodes

As your storage environment grows, you may want to consider how policy and storage pool definitions affect where workstation files are stored. Then you can define and maintain multiple storage pools in a hierarchy that allows you to contain storage costs by using sequential access storage pools in addition to disk storage pools, and provide appropriate levels of service to users.

## Estimating Space Needs in Random Access Storage Pools

To estimate the amount of storage space required for each random access (disk) storage pool:

- Decide what percentage of this storage you want to keep on disk storage space:
  - For backup storage pools, provide enough disk space to support efficient daily incremental backups.
  - For archive storage pools, provide sufficient space for a user to archive a moderate size file system without causing migration to occur.

- Establish migration thresholds to have the server migrate the remainder of the data to less expensive storage media in sequential access storage pools.  
See “Appropriate Migration Threshold Values” on page 183 for recommendations on setting migration thresholds.

### Estimating Space for Backed Up Files in a Random Access Storage Pool

To compute the total amount of space needed for all backed up files stored in a single random access (disk) storage pool, such as BACKUPOOL, use the following formula:

$$\text{backup space} = \text{avgwkstsize} * \text{utilization} * \text{versionexpansion} * \text{numwkst}$$

<b>Backup Space</b>	The total amount of storage pool disk space needed.
<b>AvgWkstSize</b>	The average data storage capacity of a workstation, in MB. For example, if the typical workstation at your installation has two 70MB hard drives, then the average workstation storage capacity is 140MB.
<b>Utilization</b>	An estimate of the fraction of each workstation disk space used, in the range 0 to 1. For example, if you expect that workstations are 75% full, then use 0.75.
<b>VersionExpansion</b>	An expansion factor (greater than 1) that takes into account the additional backup versions, as defined in the copy group. A rough estimate allows 5% additional files for each backup copy. For example, for a version limit of 2, use 1.05, and for a version limit of 3, use 1.10.
<b>NumWkst</b>	The estimated total number of workstations ADSM supports.

If compression is used, the amount of space required will be less than the total.

### Estimating Space for Archived Files in a Random Access Storage Pool

Computing the amount of storage space for archived files is more difficult, because the number of archived files generated by users is not necessarily proportional to the amount of data stored on their workstations.

To estimate the total amount of space needed for all archived files in a single random access (disk) storage pool, such as ARCHIVEPOOL, determine what percentage of user files are typically archived.

Work with policy administrators to calculate this percentage based on the number and type of archive copy groups defined. For example, if policy administrators have defined archive copy groups for only half of the policy domains in your enterprise, then you can estimate that you will need less than 50% of the amount of space you have defined for backed up files.

Because additional storage space can be added at any time, you can start with a modest amount of storage space and increase the space by adding storage volumes to the archive storage pool, as required.

## Estimating Space Needs in Sequential Access Storage Pools

To estimate the amount of space required for sequential access storage pools, consider:

- The amount of data being migrated from disk storage pools
- The length of time backed up files are retained, as defined in backup copy groups
- The length of time archived files are retained, as defined in archive copy groups
- How frequently you reclaim unused space on sequential volumes

See “Space Reclamation for Sequential Access Storage Pools” on page 190 for information about setting a reclamation threshold.

- Whether or not you use collocation to reduce the number of volume mounts required when restoring or retrieving large numbers of files from sequential volumes

If you use collocation, you may need additional tape drives.

See “Collocation on Sequential Access Storage Pools” on page 186 for information about using collocation for your storage pools.

- The type of storage devices and sequential volumes supported at your installation

---

## Defining or Updating Storage Pools

This section provides examples of how you can set up a storage pool hierarchy for an organization in your installation.

Task	Required Privilege Class
Defining storage pools	System
Update storage pool information	System or unrestricted storage

## Defining a Primary Storage Pool

When you define a primary storage pool, be prepared to provide the following information:

**Note:** Unless otherwise indicated, the following parameters apply to both random access and sequential access storage pools.

### Device class

Specifies the name of the device class assigned for the storage pool. This is a required parameter.

### Pool type

Specifies that you want to define a primary storage pool (this is the default). Updating a storage pool cannot change whether it is a primary or a copy storage pool.

### Access mode

Defines access to volumes in the storage pool for user operations (such as back up and restore) and system operations (such as reclamation and server migration). Possible values are:

**Read/Write** User and system operations can read from or write to the volumes.

**Read-Only** User operations can read from the volumes, but not write. However, system processes can move files within the volumes in the storage pool.

**Unavailable** No new writes are permitted to volumes in the storage pool from other volumes outside the storage pool. However, system processes (like reclamation) are permitted to move files within the volumes in the storage pool.

**Maximum number of scratch volumes for sequential access storage pools**

By providing a nonzero value, you specify that ADSM dynamically acquires scratch volumes.

**Maximum file size**

To exclude large files from a storage pool, set a maximum file size.

Do not set a maximum file size for the last storage pool in the hierarchy unless you want to exclude very large files from being stored in data storage.

**Migration thresholds**

Specifies a percentage of storage pool occupancy at which ADSM begins migrating files to the next storage pool (high) threshold and the percentage when migration stops (low threshold).

**Migration process**

For random access storage pools only, specifies the number of processes that are used for migrating files from this storage pool.

**Cache on random access storage pools**

Enables or disables cache. When cache is enabled, copies of migrated files are left on disk after the migration. In this way, a retrieval request can be satisfied quickly.

**Collocation for sequential access storage pools**

*Collocation* is a process in which the server attempts to keep all files belonging to a client file space on a minimal number of sequential access storage volumes.

**Reclamation threshold for sequential access storage pools**

Specifies what percentage of reclaimable space can accumulate on a volume before the server initiates a space reclamation process for the volume.

**Name of the next storage pool**

Specifies the name of the next storage pool where files can be moved.

**Reuse delay period for sequential access storage pools**

Specifies an integer that defines the number of days that must elapse after all of the files have been deleted from a volume, before the volume can be rewritten or returned to the scratch pool.

## Example: Defining A Storage Pool Hierarchy

The following steps are an example of defining a storage hierarchy.

**Defining a Disk Storage Pool:** In this example, define a disk storage pool as the destination for files backed up from the engineering department.

1. To help policy administrators identify this storage pool as a storage destination in a backup copy group used by the engineering department, name the storage pool ENGBACK1 and provide it with a meaningful description.
2. Next set a maximum file size of 5MB so that larger files are moved directly to the next storage pool in the hierarchy.
3. Then set a high migration threshold of 85% and a low migration threshold of 40% to provide sufficient free space for normal daily processing. See “Appropriate Migration Threshold Values” on page 183 for guidelines on determining migration thresholds.
4. Finally, use the default settings of read/write for the access mode and yes for cache.
5. To define the storage pool named ENGBACK1, enter:

```
define stgpool engback1 disk -  
description='disk storage pool for engineering backups' -  
maxsize=5M highmig=85 lowmig=40
```

**Defining a Tape Storage Pool:** Assume that you have already defined a device class named TAPE that describes your tape device environment.

1. The next step is to define a sequential access storage pool named ENGBACK2 and assign it to the device class named TAPE.
2. Use ENGBACK2 as the subordinate storage pool to the disk storage pool ENGBACK1. To ensure that ENGBACK2 is able to accept any size file, use the default of *no limit* for maximum file size.
3. To prevent migration from occurring from this storage pool, set the high migration threshold to 100%.
4. However, to allow the server to move large files to another storage pool if there is insufficient space on ENGBACK2, identify BACKTAPE as the next storage pool in the storage hierarchy.
5. Then to consolidate user files on separate tape volumes, set collocation to *yes* and specify that the server can request up to 100 scratch tape volumes for this storage pool.
6. Finally, use the default settings of read/write for the access mode and 60% for the reclamation threshold.
7. To define the storage pool named ENGBACK2, enter:

```
define stgpool engback2 tape -  
description='tape storage pool for engineering backups' -  
maxsize=nolimit nextstgpool=backtape highmig=100 -  
collocate=yes maxscratch=100
```



**Updating a Disk Storage Pool:** Finally, specify that ENGBACK2 is the next storage pool defined in the storage hierarchy for ENGBACK1. To update ENGBACK1, enter:

```
update stgpool engback1 nextstgpool=engback2
```

### Example: Defining A Storage Pool Hierarchy in Reverse Order

If you do not want to update each storage pool to establish the storage hierarchy, then begin by defining storage pools at the bottom of the storage hierarchy before you define the storage pools at the top of the storage hierarchy. For example, you could complete the previous three steps in two steps by entering the commands in the following order:

1. Define the bottom of the storage hierarchy first by entering:

```
define stgpool engback2 tape -  
description='tape storage pool for engineering backups' -  
maxsize=nolimit nextstgpool=backtape highmig=100 -  
collocate=yes maxscratch=100
```

2. Define the *parent* storage pool by entering:

```
define stgpool engback1 disk nextstgpool=engback2 -  
description='disk storage pool for engineering backups' -  
maxsize=360M highmig=85 lowmig=40
```

## Defining a Copy Storage Pool

When you define a copy storage pool, you should be prepared to provide the following information:

### Device class

Specifies the name of the device class assigned for the storage pool. This is a required parameter.

### Pool type

Specifies that you want to define a copy storage pool. This is a required parameter. Updating a storage pool cannot change whether it is primary or copy storage pool.

### Access mode

Defines access to volumes in the storage pool for user operations (such as back up and restore) and system operations (such as reclamation). Possible values are:

**Read/Write** User and system operations can read from or write to the volumes.

**Read-Only** User operations can read from the volumes, but not write. However, system processes can move files within the volumes in the storage pool.

**Unavailable** Specifies that users cannot access files stored on volumes in the copy storage pool. Files can be moved within the volumes of the copy storage pool, but no new writes are permitted to the volumes in the storage pool from volumes outside the storage pool.

**Maximum number of scratch volumes**

By providing a nonzero value, you specify that ADSM dynamically acquire scratch volumes.

**Collocation**

*Collocation* is a process in which the server attempts to keep all files belonging to a client node on a minimal number of sequential access storage volumes.

**Reclamation threshold**

Specifies when to initiate reclamation of volumes in the copy storage pool. Reclamation is a process that moves any remaining active, fragmented files from one volume to another volume, thus making the original volume available for reuse.

For offsite storage pool volumes, reclamation can occur regardless of whether the volume has ever been filled. An offsite volume is eligible for reclamation when the percentage of unused space on the volume is greater than the reclaim parameter value. The unused space includes both space that has never been used on the volume and space that has become empty because of file deletion.

Reclamation processing also works differently for offsite storage pool volumes. When a copy storage pool volume that is offsite becomes eligible for reclamation, the reclamation process attempts to retrieve the active files on the reclaimable volume from a primary or copy storage pool volume that is onsite, and then write these files to an available volume in the original copy storage pool.

**Reuse delay period**

Specifies an integer that defines the number of days that must elapse after all of the files have been deleted from a volume before the volume can be rewritten or returned to the scratch pool.

**Example: Defining a Copy Storage Pool**

Assume there is a need to have copies of the files stored in BACKUPPOOL, ARCHIVEPOOL, and SPACEMGPOOL (default disk storage) for disaster recovery purposes. An ADSM administrator uses the DEFINE STGPOOL command to create a copy storage pool named DISASTER-RECOVERY. It was decided to use only scratch cartridges so the maximum number of scratch volumes is set to an appropriate value.

```
define stgpool disaster-recovery cartridge pooltype=copy maxscratch=100
```

---

## Using Copy Storage Pools to Improve Data Availability

Copy storage pools enable multiple copies of files to be maintained, thus reducing the potential for data integrity loss due to media failure. If the primary file is not available or becomes corrupted, ADSM accesses and uses the duplicate file from a copy storage pool.

For example, a company has a storage hierarchy consisting of one primary random access storage pool. The files stored in the primary random access storage pool are migrated to a primary cartridge tape storage pool (device class is CARTRIDGE) called CART-POOL. Because the files are important to the function of the company, the company backs up every file in the CART-POOL.

**Note:** The company could backup the random access storage pool as well, however they have chosen to only back up the storage pool being used for migration.

The administrator decides to schedule daily incremental backups of the files in the primary storage pool. The administrator performs the following:

1. Create a copy storage pool called CART-BACKUP, with the same device class as the CART-POOL primary storage pool, by issuing the following command:

```
define stgpool cart-backup cartridge pooltype=copy
```

**Note:** All of the storage volumes in the copy storage pool CART-BACKUP are located onsite.

2. Define the same number of volumes in the copy storage pool (CART-BACKUP) as were already defined in the primary storage pool (CART-POOL) by issuing the following command:

```
define volume cart-backup <volname>
```

3. Define a schedule for backing up the primary storage pool to the copy storage pool by issuing the following command:

```
define schedule backup_cart-pool type=administrative  
cmd='backup stgpool cart-pool cart-backup'  
active=yes starttime=20:00 period=1
```

For more information about scheduling, see Chapter 7, “Scheduling Operations” on page 141.

4. Copy the files that existed in the primary storage pool (CART-POOL) prior to creating the copy storage pool (CART-BACKUP), to the copy storage pool (CART-BACKUP) by issuing the following command:

```
backup stgpool cart-pool cart-backup
```

## Recreating Damaged Files

This section explains how to recreate damaged files based on the scenario established in the previous example.

If a user tries to access a file that is stored in CART-POOL and ADSM detects a data integrity problem with the file, the server automatically retrieves the copy in CART-BACKUP and marks the primary file *damaged* so it can be recreated in the future. To recreate any *damaged* files in CART-POOL, the administrator defines a schedule that executes the following command every month:

```
restore stgpool cart-pool
```

To check for and replace any files that develop data-integrity problems in CART-POOL or in CART-BACKUP, the administrator defines schedules that issue the following commands every three months:

1. For every volume in CART-POOL the following command is executed:

```
audit volume <volname> fix=yes
```

If the AUDIT VOLUME command detects files with data-integrity errors, they are marked *damaged* and an error message is produced.

2. For every volume in CART-BACKUP the following command is executed:

```
audit volume <volname> fix=yes
```

If the AUDIT VOLUME command detects files with data-integrity errors, the copy in CART-BACKUP is deleted and a message produced.

3. Recreate *damaged* primary files by issuing the following command:

```
restore stgpool cart-pool
```

4. Produce new copies in CART-BACKUP by issuing the following command:

```
backup stgpool cart-pool cart-backup
```

---

## Backing Up Storage Pools

Administrators can back up primary storage pools into copy storage pools.

Task	Required Privilege Class
Back up storage pools	System, unrestricted storage, or restricted storage for the copy storage pool

The BACKUP STGPOOL command is used to copy files into a copy storage pool. Because the copies are made incrementally, the backup process may run as long as required to back up the primary storage pool or be cancelled if desired. Reissuing the BACKUP STGPOOL command allows the backup to continue from the spot the backup was cancelled. For example, to back up the ARCHIVEPOOL primary pool to the RECOVERYPOOL copy pool, enter:

```
backup stgpool archivepool recoverypool
```

The BACKUP STGPOOL command can also be scheduled. The administrator can define schedules to initiate incremental backups of files in the primary storage pools. For example, to back up the BACKUPPOOL, ARCHIVEPOOL, and the TAPEPOOL every night, the following commands are scheduled:

```
backup stgpool backuppool disaster-recovery maxprocess=4  
backup stgpool archivepool disaster-recovery maxprocess=4  
backup stgpool tapepool disaster-recovery maxprocess=4
```

These commands use multiple, parallel processes to perform an incremental backup of each primary storage pool to the copy pool. The only files backed up to the DISASTER-RECOVERY pool are files for which a copy does not already exist in the copy storage pool. See Chapter 7, “Scheduling Operations” on page 141 for information about scheduling commands.

**Notes:**

1. Backing up storage pools places additional space requirements on the database.
2. If a copy is to be generated in a specific copy storage pool and a copy already exists with the same insertion date, no action is taken.
3. File copies stored in a copy storage pool do not migrate.
4. Cached files are not backed up.
5. When setting the MAXPROCESS parameter, consideration should be given to the number of mount points and drives that can be dedicated to this operation.

See “Using Storage Pool Backup Features” on page 280 for more information about using storage pool back up in your disaster recovery strategy.

## Full Storage Pool Backups

Incremental backups of storage pools, along with database backups and periodic reclamation of offsite volumes, should provide protection to meet most recovery requirements. However, you may still wish to perform occasional full backups of primary storage pools. Reasons for this may include:

- Concern for shelf life of the media being used for backups
- Legal or audit requirements that specify how backups should be performed
- Familiarity with a full backup and applied incremental approach

A full storage pool backup can be achieved at any time by backing up to a new copy storage pool. Further backups to this new copy storage pool will be incremental.

### **Example: Full Storage Pool Backup**

Suppose you wish to take weekly full backups and daily incrementals and you want to keep 4 weeks worth of backups. This can be accomplished by backing up to a new copy storage pool each week. When a copy storage pool is 4 weeks old, you can delete all the volumes in that copy storage pool (using the DELETE VOLUME command with DISCARDATA=YES). The copy storage pool and its volumes could then be reused for the new weeks's backups.

There are some drawbacks to this approach that must be understood:

- There will be a lot of database activity introduced with deleting the volumes regularly.
- This approach will maintain 5 copies of each file, including the primary copy .
- Information is kept in the database for each copy of a file in a copy storage pool.

**Note:** This information is not a full database entry such as is used for the primary copy, but it will still take up database space.

---

## **Monitoring the Use of Storage Pools**

Any administrator can query for information about a storage pool by viewing a standard or a detailed report. Use these reports to monitor storage pool usage, including:

- The use of space in your disk and sequential access storage pools
- Migration of data from one to storage pool to the next storage pool in the storage hierarchy
- The use of disk space by cached copies of migrated files

## **Monitoring the Use of Storage Pool Space**

To query the server to view a standard report for all storage pools defined to the system, enter:

```
query stgpool
```

Figure 43 on page 203 shows a standard report with all storage pools defined to the system. To monitor the use of storage pool space, review the *Estimated Capacity* and *%Util* columns.

Storage Pool Name	Device Class Name	Estimated Capacity (MB)	%Util	%Migr	High Mig%	Low Mig%	Next Storage Pool
ARCHIVEPOOL	DISK	0.0	0.0	0.0	90	70	
BACKTAPE	TAPE	180.0	85.0	100.0	90	70	
BACKUPPOOL	DISK	80.0	51.6	51.6	50	30	BACKTAPE
COPYPOOL	TAPE	300.0	42.0				
ENGBACK1	DISK	0.0	0.0	0.0	90	70	
ENGBACK2	DISK	0.0	0.0	0.0	90	70	

Figure 43. Information about Storage Pools

*Estimated capacity* specifies the available space of the storage pool in megabytes.

For disk storage pools, this value reflects the total amount of available space in the storage pool, including any volumes that are varied offline.

For sequential access storage pools, the estimated capacity value is an estimate of the total amount of available space on all volumes in the storage pool, including volumes that have *unavailable*, *read only*, *offsite*, or *destroyed* access mode and all scratch volumes that can be acquired in this storage pool. Recall that volumes in a sequential access storage pool, unlike those in a disk storage pool, do not contain preallocated space. Rather, data is written to these volumes as necessary until the end of the volume is reached. It is for this reason that the estimated capacity is truly an *estimate* of the amount of available space in a sequential access storage pool.

*%Util* specifies, as a percentage, the use of each storage pool.

For disk storage pools, this value reflects the total number of disk blocks currently allocated to ADSM. Space is allocated for backed up, archived or client-migrated files that are eligible for server migration, cached files which are copies of server-migrated files, and files that reside on any volumes which are varied offline.

**Note:** The value for %UTIL can be slightly higher than the value for %MIGR if you query for storage pool information while a backup or archive transaction is in progress. The value for %UTIL is determined by the amount of space actually allocated (while the transaction is in progress), while the value for %MIGR only represents the space occupied by *committed* files. At the end of the transaction, %UTIL and %MIGR become synchronized.

For sequential access storage pools, this value is the percentage of the total bytes of storage available that are currently being used to store active (non-expired) data. Because the server can only estimate the available capacity of a sequential access storage pool, this percentage also reflects an estimate of the actual utilization of the storage pool.

### Example: Monitoring the Capacity of a Backup Storage Pool

Figure 43 shows that the estimated capacity for a disk storage pool named BACKUPPOOL is 80MB, which is the amount of available space on disk storage. More than half (51.6%) of the available space is occupied by either backup files or cached copies of backup files.

The estimated capacity for the tape storage pool named BACKTAPE is 180MB, which is the total estimated space available on all tape volumes in the storage pool. This report shows that 85% of the estimated space is currently being used to store workstation files.

**Note:** This report also shows that volumes have not yet been defined to either the ENGBACK1 or ENGBACK2 storage pools, since both storage pools show an estimated capacity of 0.0MB.

## Monitoring Migration Thresholds

Four fields on the standard storage pool report provide you with information about the migration process. They include:

### **%Migr**

Specifies the percentage of data in each storage pool that can be migrated. This value is used to determine when to start or stop migration.

For disk storage pools, this value represents the amount of disk space occupied by backed up, archived, or client-migrated files that can be migrated to another storage pool, including files on volumes that are varied offline. Cached data are excluded in the %MIGR value.

For sequential access storage pools, this value is the percentage of the total volumes in the storage pool that actually contain data at the moment. For example, assume a storage pool has four explicitly defined volumes, and a maximum scratch value of six volumes. If two volumes are actually contain data at the moment, then %Migr will be 20% (this field is left blank for copy storage pools).

### **High Migr%**

Specifies when ADSM can begin migrating data from this storage pool. Migration can begin when the percentage of data that can be migrated reaches this threshold (this field is left blank for copy storage pools).

### **Low Migr%**

Specifies when ADSM can stop migrating data from this storage pool. Migration can end when the percentage of data that can be migrated falls below this threshold (this field is left blank for copy storage pools).

### **Next Storage Pool**

Specifies the primary storage pool destination to which data is migrated (this field is left blank for copy storage pools).

## **Example: Monitoring the Migration of Data Between Storage Pools**

ADSM sets a default of 90% for the high migration threshold and 70% for the low migration threshold for each primary storage pool.

Figure 43 on page 203 shows that the predefined migration thresholds for BACKUPPOOL storage pool have been updated to 50% for the *high migration threshold* and 30% for the *low migration threshold*.

When the amount of data stored in the storage pool reaches 50%, the server can begin to migrate files to BACKTAPE.

To monitor the migration of files from BACKUPPOOL to BACKTAPE, enter:



```
query stgpool back*
```

If caching is on for a disk storage and files are migrated, the %UTIL value does not change since the cached files still occupy space in the disk pool. However, the %MIGR value decreases since this space is no longer migratable because migration has already occurred. See Figure 44 for an example.

Storage Pool Name	Device Class Name	Estimated Capacity (MB)	%Util	%Migr	High Mig%	Low Mig%	Next Storage Pool
BACKTAPE	TAPE	180.0	95.2	100.0	90	70	
BACKUPPOOL	DISK	80.0	51.6	28.8	50	30	BACKTAPE

Figure 44. Information on Backup Storage Pools

At this point, a system administrator can:

- Cancel the migration process  
See “Canceling the Migration Process” on page 206 for additional information.
- End the migration process  
See “Ending the Migration Process” on page 206 for additional information.
- Provide additional space  
See “Providing Additional Space for the Migration Process” on page 206 for additional information.

You can query the server to monitor the migration process by entering:

```
query process
```

A message similar to Figure 45 is displayed:

Process Number	Process Description	Status
2	Migration	Disk Storage Pool BACKUPPOOL, Moved Files: 1086, Moved Bytes: 25555579, Unreadable Files: 0, Unreadable Bytes: 0

Figure 45. Information on the Migration Process

When migration is finished, the server displays the following message:

```
ANR1101I Migration ended for storage pool BACKUPPOOL.
```

## Canceling the Migration Process

Before a system administrator can cancel the migration process, determine the identification number of the background migration process by entering:

```
query process
```

A message similar to Figure 46 is displayed:

Process Number	Process Description	Status
1	Migration	ANR1113W Migration suspended for storage pool BACKUPPOOL - insufficient space in subordinate storage pool.

Figure 46. Number of the Migration Process

Then a system administrator can cancel the migration process by entering:

```
cancel process 1
```

## Ending the Migration Process

Task	Required Privilege Class
Update a storage pool	System or unrestricted storage

You can update the storage pool to cause an immediate end to the migration process. Some errors cause the server to continue attempting to restart the migration process after 60 seconds. If the problem still exists after several minutes, the migration process will end. Depending on your environment, you can:

- Set higher migration thresholds for the disk storage pool to delay the server from initiating migration
- Add volumes to the disk storage pool to increase the storage capacity of BACKUPPOOL, thereby decreasing the migration percentage (%Migr)

**Note:** This would only be done if you received an out of space message.

## Providing Additional Space for the Migration Process

Task	Required Privilege Class
Update a storage pool	System or unrestricted storage

You can update the storage pool to provide additional storage volumes to complete the migration process. Add volumes to the tape storage pool or increase the maximum number of scratch tapes to increase the storage capacity of BACKTAPE. The server attempts to restart the migration process every 60 seconds for several minutes and then will terminate the migration process.

## Monitoring the Use of Cache Space on Disk Storage

The %UTIL value includes cached data on a volume (when cache is enabled) and the %MIGR value excludes cached data. Therefore, when cache is enabled and migration occurs, the %MIGR value decreases while the %UTIL value remains the same. The %UTIL value remains the same because the migrated data remains on the volume as cached data. In this case, the %UTIL value only decreases when the cached data expires.

If you update a storage pool from CACHE=YES to CACHE=NO, the cached files will not disappear immediately. The %UTIL value will be unchanged. The cache space will be reclaimed over time as the server needs the space, and no additional cached files will be created.

To determine whether cache is being used on disk storage and to monitor how much space is being used by cached copies, query the server for a detailed storage pool report. For example, to request a detailed report for BACKUPPOOL, enter:

```
query stgpool backuppool format=detailed
```

Figure 47 displays a detailed report for the storage pool.

```
Storage Pool Name: BACKUPPOOL
Storage Pool Type: PRIMARY
Device Class Name: DISK
Estimated Capacity (MB): 80.0
    %Util: 42.0
    %Migr: 29.6
    High Mig%: 50
    Low Mig%: 30
Migration Processes:
    Next Storage Pool: BACKTAPE
Maximum Size Threshold: No Limit
    Access: Read/Write
Description:
    Cache Migrated Files?: Yes
    Collocate?:
Reclamation Threshold:
Maximum Scratch Volumes Allowed:
    Delay Period for Volume Reuse: 0 Day(s)
    Migration in Progress?: Yes
    Amount Migrated (MB): 0.10
Elapsed Migration Time (seconds): 5
    Reclamation in Progress?:
Volume Being Migrated/Reclaimed:
    Last Update by (administrator): SERVER_CONSOLE
    Last Update Date/Time: 04/07/1995 16:47:49
```

Figure 47. Detailed Storage Pool Report

When *Cache Migrated Files?* is set to *yes*, the value for %Util should not change because of migration, because cached copies of migrated files remain in disk storage.

This example shows that utilization remains at 42%, even after files have been migrated to the BACKTAPE storage pool, and the current amount of data eligible for migration is 29.6%.

When *Cache Migrated Files?* is set to *no*, the value for %Util more closely matches the value for %Migr because cached copies are not retained in disk storage.

## Requesting Information on Storage Occupancy

Task	Required Privilege Class
Query the server for information about data storage	Any administrator

Any administrator can request information about data storage occupancy. Use the QUERY OCCUPANCY command for reports with information broken out by node or file space. Use the QUERY STGPOOL command for any other information. Use this report to determine the amount of space used by:

- Client node and file space
- Storage pool or device class
- Type (backup, archive, or space managed)

You can also use this report to evaluate the average size of workstation files stored in data storage.

### Amount of Space Used by Client Node

Any administrator can request information:

- About how much data has been backed up or archived to the server by file space
- About the amount of storage space being used by client node and file space

For information about querying the server for file space information, see “Requesting File Space Information” on page 108.

To determine the amount of server storage space used by the /home file space belonging to the client node SSTEINER, for example, enter:

```
query occupancy ssteiner /home
```

Remember that file space names are case-sensitive and must be entered using the exact capitalization used when the file space name was initially defined by the associated ADSM client system. Use the QUERY FILESPACE command to determine the correct capitalization.

Figure 48 on page 209 shows the number of files backed up or archived from the /home file space belonging to SSTEINER. The report also shows how much space is occupied in each storage pool.

Node Name	Filespace Name	Storage Pool Name	Number of Files	Space Occupied (MB)
----- SSTEINER	----- /home	----- ENGBACK1	----- 513	----- 3.52

Figure 48. A Report of the Occupancy of Storage Pools by Client Node

### Amount of Space Used by Storage Pool or Device Class

You can monitor the amount of space being used by an individual storage pool, a group of storage pools, or storage pools categorized by a particular device class. Creating occupancy reports on a regular basis can help you with capacity planning.

For example, to query the server for the amount of data stored in backup tape storage pools belonging to the CARTRIDGE device class, enter:

```
query occupancy devclass=cartridge
```

Figure 49 displays a report on the occupancy of tape storage pools assigned to the CARTRIDGE device class.

Node Name	Filespace Name	Storage Pool Name	Number of Files	Space Occupied (MB)
----- HTANG	----- OS2C	----- ARCTAPE	----- 5	----- .92
HTANG	OS2C	BACKTAPE	21	1.02
PEASE	/home/peas- e/dir	ARCTAPE	492	18.40
PEASE	/home/peas- e/dir	BACKTAPE	33	7.60
PEASE	/home/peas- e/dir1	BACKTAPE	2	.80
TOMC	/home/tomc /driver5	ARCTAPE	573	20.85
TOMC	/home	BACKTAPE	13	2.02

Figure 49. A Report on the Occupancy of Storage Pools by Device Class

### Amount of Space Used by Backed Up, Archived or Migrated Files

Finally, you can query the server for the amount of space used by backed up and archived files, and files migrated from client nodes. By determining the average size of workstation files stored in data storage, you can estimate how much storage capacity you might need when registering new client nodes to the server. See “Estimating Space Needs for Storage Pools” on page 192 and “Estimating Space for Archived Files in a Random Access Storage Pool” on page 193 for information about planning storage space.

To request a report about backup versions stored in the disk storage pool named BACKUPPOOL, for example, enter:

```
query occupancy stgpool=backuppool type=backup
```

Figure 50 displays a report on the amount of data storage used for backed up files.

Node Name	Filespace Name	Storage Pool Name	Number of Files	Space Occupied (MB)
HTANG	OS2C	BACKUPPOOL	513	23.52
HTANG	OS2D	BACKUPPOOL	573	20.85
PEASE	/marketing	BACKUPPOOL	132	12.90
PEASE	/business	BACKUPPOOL	365	13.68
TOMC	/	BACKUPPOOL	177	21.27

Figure 50. A Report of the Occupancy of Backed Up Files in Storage Pools

To determine the average size of backup versions stored in BACKUPPOOL, complete the following steps using the data provided in Figure 50:

1. Add the number of megabytes of space occupied by backup versions.  
In this example, backup versions occupy 92.22MB of space in BACKUPPOOL.
2. Add the number of files stored in the storage pool.  
In this example, 1760 backup versions reside in BACKUPPOOL.
3. Divide the space occupied by the number of files to determine the average size of each file backed up to the BACKUPPOOL.  
In this example, the average size of each workstation file backed up to BACKUPPOOL is about 0.05MB, or approximately 50KB.

You can use this average to estimate the capacity required for additional storage pools that are defined to ADSM.

## Deleting a Storage Pool

Task	Required Privilege Class
Delete storage pools	System

Before a storage pool can be deleted, ensure that:

- All volumes within the storage pool have been deleted

Ensure that you have saved any readable data that you want to preserve by issuing the MOVE DATA command. Moving all of the data that you want to preserve may require you to issue the MOVE DATA command several times.

Before you begin deleting all volumes that belong to the storage pool, change the access mode of the storage pool to unavailable so that no files can be written to or read from volumes in the storage pool.

See “Deleting a Storage Pool Volume with Data” on page 229 for information about deleting storage volumes.

- The storage pool is not identified as the next storage pool within the storage hierarchy

To determine whether this storage pool is referenced as the next storage pool within the storage hierarchy, query for storage pool information as described in “Monitoring the Use of Storage Pool Space” on page 202.

Update any storage pool definitions to remove this storage pool as a subordinate storage pool in the storage hierarchy by performing one of the following:

- Naming another storage pool as the next storage pool in the storage hierarchy
- Entering double quotes ( ' ' ) on the *next* parameter to remove this storage pool from the storage hierarchy definition.

Make sure that the storage pool to be deleted is not specified as the destination for any copy groups in any management classes within the active policy set of any domains. If this pool is a destination and the pool is deleted, backup/archive operations fail because there is no storage space to store the data. Also, a storage pool to be deleted cannot be the destination for space managed files if the destination is specified in any management classes.

See “Defining or Updating Storage Pools” on page 194 for information about updating storage pool definitions.

---

## Restoring Storage Pools

An administrator can recreate files in a primary storage pool using duplicate copies in a copy storage pool by issuing the RESTORE STGPOOL command.

Task	Required Privilege Class
Restoring storage pools	System, unrestricted storage, or restricted storage

Use the RESTORE STGPOOL command to restore files from one or more copy storage pools to a primary storage pool. The files must have been copied to the copy storage pools by using the BACKUP STGPOOL command.

The RESTORE STGPOOL command restores specified primary storage pools that have files with the following problems:

- The primary copy of the file has been identified as having data-integrity errors during a previous operation.

**Note:** Files with data-integrity errors are marked as damaged.

- The primary copy of the file resides on a volume which has an access mode of DESTROYED.

The RESTORE STGPOOL command with the PREVIEW=YES parameter can be used to identify volumes that contain damaged primary files. During restore processing, a message is issued for every volume in the restored storage pool that contains damaged, non-cached files. Use the QUERY CONTENT command to identify damaged, primary files on a specific volume.

**Note:** Cached copies of files are never restored. Any cached files which have been identified as having data-integrity errors or which reside on a *destroyed* volume will be removed from the database during restore processing.

After the files are restored, the old references to these files in the primary storage pool are deleted from the database. This means that ADSM will now locate these files on the volumes to which they were restored, rather than on the volumes on which they were previously stored. If a destroyed volume becomes empty because all files have been restored to other locations, the destroyed volume is automatically deleted from the database.

If the backup file copies are moved or deleted during restore processing, the restore may be incomplete. Therefore, do not issue the following commands for copy storage pool volumes while restore processing is in progress:

- MOVE DATA
- DELETE VOLUME (DISCARDDATA=YES)
- AUDIT VOLUME (FIX=YES)

In addition, you can delay reclamation processing for your copy storage pools by setting the RECLAIM parameter to 100 (using the UPDATE STGPOOL command) while restore processing is in progress.

**Note:** Copy storage pool volumes that are located offsite must be returned to the onsite location for restore processing.

The RESTORE STGPOOL command generates a background process that can be canceled with the CANCEL PROCESS command. If a RESTORE STGPOOL background process is canceled, some files may have already been restored prior to the cancellation. To display information about background processes, use the QUERY PROCESS command.

When you restore a storage pool, be prepared to provide the following information:

**Primary storage pool**

Specifies the name of the primary storage pool that is being restored.

**Copy storage pool**

Specifies the name of the copy storage pool from which the files are to be restored.

**New storage pool**

Specifies the name of the new primary storage pool to which to restore the files.

**Maximum number of processes**

Specifies the number of parallel processes that are used for restoring files.

**Preview**

Specifies whether you want to preview the restore operation before it is actually performed.

See “Recovering by Using Backed Up Copies of Storage Pools” on page 280 for an example of using the RESTORE STGPOOL command.



---

## Chapter 11. Managing Storage Pool Volumes

Task	Required Privilege Class
Define volumes in any storage pool	System or unrestricted storage
Define volumes in specific storage pools	System, unrestricted storage, or restricted storage for those pools
Update volumes	System or operator
Display information about volumes	Any administrator

The sections listed in the following table begin at the indicated pages.

Section	Page
<b>Concepts:</b>	
Storage pool volumes	213
Access mode for storage pool volumes	214
<b>Tasks:</b>	
Adding random access storage pool volumes	215
Defining storage pool volumes	216
Updating storage pool volumes	216
Monitoring the use of storage pool volumes	217
Managing data storage	219
Moving files from one volume to another volume	225
Deleting storage pool volumes	228
Restoring storage pool volumes	230

Most tasks presented in this chapter can be performed by using either the graphical user interface or the command-line interface. Table 9 on page 43 shows whether a task can be performed on the graphical user interface, the command line-interface, or both.

For information about issuing ADSM commands, refer to *ADSM Administrator's Reference*. For help performing a task on the graphical user interface, see the procedure described in "Using Online Help" on page 15.

---

### Storage Pool Volumes

Volumes in storage pools contain backed up, archived, or migrated data from clients. Storage pools are either random access or sequential access, depending on the type of volume assigned to the pool.

Random access storage pools are always associated with the DISK device class and all volumes in this type of storage pool have the same form.

Each volume defined in a sequential access storage pool must be the same type as the device type of the associated device class. The device types are:

<b>VSAM linear data set</b>	A collection of logically related data records stored on DASD or tape.
<b>Tape cartridge</b>	Sequential storage media containing magnetic tape in a protective, square housing.

---

## Access Mode for Storage Pool Volumes

Access to any volume in a storage pool is determined by the access mode assigned to that volume:

<b>Read/write</b>	<p>Allows files to be read from or written to a volume in the storage pool.</p> <p>If the server is cannot write to a read/write access volume, the server automatically changes the access mode to read-only.</p>
<b>Read-only</b>	Allows files to be read from but not written to a disk or tape volume.
<b>Unavailable</b>	Specifies that the volume is not available for any type of access by the ADSM server.
<b>Destroyed</b>	<p>Specifies that a primary storage pool volume has been permanently damaged. Neither users nor system processes (like migration) can access files stored on the volume.</p> <p>Only volumes in primary storage pools can be updated to destroyed.</p> <p>If you update a random access storage pool to destroyed, you cannot vary the volume online. If you update a sequential access storage pool volume to destroyed, ADSM does not attempt to mount the volume.</p> <p>This access mode is used to indicate an entire volume which should be restored using the RESTORE STGPOOL or RESTORE VOLUME command. After all files on a destroyed volume are restored, the volume is automatically deleted from the database.</p> <p>If a volume contains no files and the UPDATE VOLUME command is used to change the access mode to destroyed, the volume is deleted from the database.</p>
<b>Offsite</b>	<p>Specifies that a copy storage pool volume is at an offsite location and therefore cannot be mounted. Use this mode to designate offsite volumes so that:</p> <ul style="list-style-type: none"> <li>• You will know which volumes are located offsite</li> <li>• Mount requests will not be generated for offsite volumes</li> <li>• Data can be reclaimed or moved from offsite volumes by retrieving files from other storage pools</li> <li>• Empty, offsite scratch volumes that are not deleted from the copy storage pool</li> </ul> <p>Only volumes in a copy storage pool can be updated to offsite.</p>

---

## Adding Random Access Storage Pool Volumes

When you define a random access storage pool volume, it must occupy at least 1MB of space.

Create additional storage pool volumes by performing the following steps:

1. Allocate and format the new volume by running one of the following two jobs.:

If the SMS address space is *not* active, edit and run the ANRFMT job from the ADSM.SAMPLIB library

The parameters you must edit before running this job are as follows:

- a. The job card parameters. Customize them for your installation.
- b. The NAME, LINEAR MB, and VOL parameters in the DEFINE CLUSTER statement. For example, if the name of your first new storage pool volume is adsm.storage.pool003, the size is 100 MB, and the volume is ADSM01, edit the DEFINE CLUSTER statement as follows:

```
DEFINE CLUSTER( NAME(adsm.storage.pool003) LINEAR MB(100) VOL(ADSM01) )
```

- c. The PARM parameter in the ADSMFMT EXEC statement. For example, enter the name of the new recovery log volume as follows:

```
PARM='/adsm.storage.pool003'
```

For more information about the ANRFMT job, see “Using the ANRFMT Job” on page 71.

**OR**

If the SMS address space is active and you have installed the ANRFMT2 job in a procedure library, enter the following command at the system console:

```
s anrfmt,name=adsm.storage.pool003,mb=100,volser=ADSM01
```

For more information about the ANRFMT2 job, see “Using the ANRFMT2 Job.” on page 73.

2. Edit and run the ANRFMT or ANRFMT2 job again to allocate and format the second new storage pool volume as you did in the previous step. For example, if the name of your second new storage pool is adsm.storage.pool004, enter this name in the appropriate parameter fields.
3. Assign the newly created storage pool volumes to the storage pools by entering the following commands:

```
define volume backuppool adsm.storage.pool003
define volume archivepool adsm.storage.pool004
```

See “Step 1: Defining Disk Volumes” on page 75 for details about defining volumes to ADSM.

4. If you are running HSM feature on the ADSM server, you can also allocate and format additional volumes for the space management storage pool by using the ANRFMT or ANRFMT2 jobs. After you allocate and format the new volume, you can assign it to the space management storage pool.

For example, if you named the new volume adsm.storage.pool005, enter the following command to assign it to the space management storage pool:

```
define volume spacemgpool adsm.storage.pool005
```

---

## Defining Storage Pool Volumes

The ADSM server can use dynamically acquired scratch volumes, predefined volumes, or a combination in a sequential access storage pool. Volumes in a random access storage pool must be predefined.

When you define a storage pool volume, you inform that server that the volume is available for use when storing backup, archive, or space-managed data. Before a scratch tape volume can be used, it must have a standard label.

To define a volume named VOL1 in the ENGBACK3 storage pool, enter:

```
define volume engback3 vol1
```

**Note on Sequential Access Volumes:** You do not have to define volumes in sequential storage pools if you use the MAXSCRATCH parameter when you define or update the storage pool. Setting MAXSCRATCH to a nonzero value lets the storage pool dynamically acquire volumes as needed. The volumes are automatically defined as they are acquired; they are also automatically deleted from the storage pool when the server no longer needs them.

---

## Updating Storage Pool Volumes

Use the UPDATE VOLUME command to update the attributes of an existing random or sequential access storage pool volume. This command can be used to update volumes assigned to primary or copy storage pools.

**Note:** You can reset any error state associated with a volume by updating the volume to an access mode of READWRITE.

A random access volume must be varied offline before you can update the volume to *unavailable* or *destroyed*. To vary a volume offline, use the VARY command.

If a scratch volume with a status of EMPTY and an access mode of OFFSITE is updated so that the access mode is READWRITE, READONLY, or UNAVAILABLE, the volume will be deleted from the database.

When using the UPDATE VOLUME command, be prepared to supply the following information:

**Volume name**

Specifies the name of the storage pool volume to be updated. This parameter is optional.

**Access**

Specifies how users and system processes (like migration) can access files in the storage pool volume.

**Location**

Specifies the location of the volume. This parameter is optional. It can be specified only for volumes in sequential-access storage pools.

**Storage pool**

Specifies the name of the storage pool for volumes to be updated.

**Device class**

Specifies the name of the device class for volumes to be updated. This parameter can be used to restrict the update by device class.

**Access**

Specifies the current access mode of volumes to be updated.

**Status**

Specifies the status of volumes to be updated.

**Preview**

Specifies whether you want to preview the update operation before it is actually performed.

An example of when to use the UPDATE VOLUME command might be if you accidentally damage VOL1, you can change the access mode to unavailable so that no data can be written to or read from the volume. Enter the following:

```
update volume vol1 access=unavailable
```

---

## Monitoring the Use of Storage Pool Volumes

You can request that the server produce a report about storage pool volumes. You can query the server for general information about storage pool volumes, or you can view a detailed report to evaluate:

- Current access mode and status of the volume
- Amount of available space on the volume
- Amount of reclaimable space on a sequential access volume
- Location

## Requesting General Information about Storage Pool Volumes

To query the server for general information about all volumes defined to the server, enter:

```
query volume
```

Volume Name	Storage Pool Name	Device Class	Estimated Capacity (MB)	%Util	Volume Status
ATF001	BACKUPPOOL	DISK	400.1	18.3	On-Line
DA0001	ARCTAPE	CARTRIDGE	0.0	0.0	Empty
DA0002	ARCTAPE	CARTRIDGE	0.0	0.0	Empty
DB0001	BACKTAPE	CARTRIDGE	0.0	0.0	Empty
DB0001	BACKTAPE	CARTRIDGE	180.0	11.7	Filling
DB0002	BACKTAPE	CARTRIDGE	0.0	0.0	Empty

Figure 51. Standard Information About Storage Pool Volumes

## Requesting Detailed Information about Storage Pool Volumes

To query the server for a detailed report on the volume named DB0001 in the storage pool named BACKTAPE, enter:

```
query volume db0001 stgpool=backtape format=detailed
```

Figure 52 shows the output of this detailed query.

```
Volume Name: DB0001
Storage Pool Name: BACKTAPE
Device Class: CARTRIDGE
Estimated Capacity (MB): 180.0
%Util: 11.7
Volume Status: Filling
Access: Read/Write
Pct. Reclaimable Space: 1
Scratch Volume?: No
In Error State?: No
Number of Writable Sides: 1
Number of Times Mounted: 5
Write Pass Number: 2
Approx. Date Last Written: 03/29/1993 12:04:30
Approx. Date Last Read: 03/30/1993 09:30:02
Number of Write Errors: 0
Number of Read Errors: 0
Last Update by (administrator): PENNER
Last Update Date/Time: 03/23/1993 13:16:04
```

Figure 52. Detailed Information about a Storage Volume

Use this report to:

- Ensure that the volume is available for use.

Check the *volume status* to see if a disk volume has been varied offline, or if a sequential access volume is currently being filled with data.

Check the *access mode* to determine whether files can be read from or written to this volume.

- Monitor the use of storage space.

The *estimated capacity* is determined by the device class associated with the storage pool to which this volume belongs. Based on the estimated capacity, the system tracks the percentage of space occupied by client files. In this example, 11.7% of the estimated capacity is currently in use.

- Monitor the life of a sequential access volume.

In this example, DB0001 is not a scratch volume, which means that it will be reused by the BACKTAPE storage pool after space has been reclaimed or deleted from the volume.

The *write pass number* indicates the number of times the volume has been written to, starting from the beginning of the volume. A value of one indicates that a volume is being used for the first time. In this example, DB0001 has a write pass number of two, which indicates space on this volume may have been reclaimed or deleted once before. Be sure to compare this value to the specifications provided with the media that you are using. In particular, the manufacturer recommendations for the maximum number of write passes for some types of tape media may require that you retire your tape volumes after reaching the limit in order to ensure the integrity of your data.

Use the *number of times mounted* and the *approximate date last written to or read from* to help you estimate the life of the volume. For example, if more than six months have passed since the last time this volume has been written to or read from, you should audit the volume to ensure that files can still be accessed. See “Auditing a Storage Pool Volume” on page 220 for information about auditing a volume.

- Monitor the error status of the volume.

The server reports when the volume is in an error state and automatically updates the access mode of the volume to read-only. The *number of write errors* and *number of read errors* indicate the type and severity of the problem. Audit a volume when it is placed in error state. See “Auditing a Storage Pool Volume” on page 220 for information about auditing a volume.

- Determine the location of an offsite volume.

The location of a volume in a sequential-access storage pool is shown if the optional LOCATION parameter was used when the volume was defined or during a previous UPDATE VOLUME command.

---

## Managing Data Storage

Task	Required Privilege Class
Audit or move files from volumes in storage pools over which they have authority	Restricted storage privilege
Audit or move files from a volume in any storage pool	System privilege, unrestricted storage privilege
Display information about files in storage pool volumes	Any administrator

Use this section to help you:

- Audit storage pool volumes for data integrity
- Determine what is contained on a storage pool volume
- Move files from one volume to other available volumes

## Auditing a Storage Pool Volume

When users back up, archive, or migrate files, the server retains information about the files in the database, while actually storing backup versions and client-migrated and archive copies in data storage. If there are inconsistencies between the information in the database and the files in a storage pool volume, users cannot restore or retrieve the files from data storage.

For example, you experience a disk failure on a volume in a random access storage pool. The server has already recorded in the database that files were stored on the failed disk. However, because of the disk failure, the backed up files are irretrievable.

To ensure that all recorded files are accessible on volumes in a storage pool, you audit any failed volumes by using the `AUDIT VOLUME` command. An audit checks for any logical inconsistencies between the database and the storage pool volume. You can determine what volumes belong to a storage pool by querying the server for volume information as described in “Requesting General Information about Storage Pool Volumes” on page 218. If files with integrity errors are detected, the handling of these files depends on the following:

- The type of storage pool where the volume belongs
- The `FIX` option of the `AUDIT VOLUME` command
- The location of file copies

For a volume in a primary storage pool, the values for the `FIX` parameter on an `AUDIT VOLUME` command have the following effects:

### **Fix=No**

ADSM reports, but does not delete, any database records that refer to files found with logical inconsistencies.

If the `AUDIT VOLUME` command detects a data-integrity error in a file:

- ADSM marks the file as *damaged* in the database. If a backup copy is stored in a copy storage pool, the file can be restored using the `RESTORE VOLUME` or `RESTORE STGPOOL` command.
- If the file is a cached copy, references to the file on this volume can be deleted by issuing the `AUDIT VOLUME` command and specifying `FIX=YES`.

If the `AUDIT VOLUME` command does not detect a data-integrity error in a file that had previously been marked as damaged, the state of the file is reset so that the file can be used. This provides a means for resetting the state of damaged files if it is determined that the errors were caused by correctable hardware problems such as a dirty tape head.

### **Fix=Yes**

ADSM fixes any inconsistencies as they are detected.

If the `AUDIT VOLUME` command detects a data-integrity error in a file:



- If a backup copy is not stored in a copy storage pool, ADSM deletes all database records that refer to the file.
- If a backup copy is stored in a copy storage pool, ADSM marks the file as damaged in the database. The file can then be restored using the RESTORE VOLUME or RESTORE STGPOOL command.
- If the file is a cached copy, ADSM deletes the database records that refer to the cached file. The primary file is stored on another volume.

If the AUDIT VOLUME command does not detect a data-integrity error in a file that had previously been marked as damaged, ADSM resets the state of the file so that it can be used. This provides a means for resetting the state of damaged files if it is determined that the errors were caused by correctable hardware problems such as a dirty tape head.

For a volume in a copy storage pool, the values for the FIX parameter on an AUDIT VOLUME command have the following effects:

**FIX=NO**

The error is reported and the file copy marked as damaged in the database.

**FIX=YES**

References to the file on the audited volume are deleted.

**Considerations for Auditing Storage Pool Volumes**

When you audit a volume, the server initiates a background process that sends progress messages to the server console and activity log. During the audit process, ADSM checks for any inconsistencies between information about files stored in the database and the actual files stored on a volume in a storage pool.

You should audit a volume when:

- The volume is damaged
- The volume has not been accessed for a long period of time, for example, after six months
- A read or write error occurs while accessing the volume
- The database has been restored to an earlier point in time, and the volume is either a disk volume or a volume that was identified as being reused or deleted since the database backup took place

When you audit a volume, a background process is started as the server checks for inconsistencies between the contents of the database and the contents of the volume. During the auditing process, the server:

- Records processing information in the activity log
- Sends informational messages about processing to the server console
- Prevents new files from being written to the volume

You can specify whether you want the server to correct the database in the event that inconsistencies are detected. The system default is to report inconsistencies that are found, but to not correct the errors.

## Auditing a Volume in a Disk Storage Pool

To audit a disk volume named ADSM.STORAGE.POOL001, for example, and have only summary messages sent to the activity log and server console, enter:

```
audit volume adsm.storage.pool001 quiet=yes
```

The audit volume process is run in the background, and the server returns an informational message as follows:

```
ANR2313I Audit Volume NOFIX process started for volume  
ADSM.STORAGE.POOL001 (process id 4).
```

To view the status of the audit volume process, enter:

```
query process
```

The following figure displays an example of the report you receive about the audit volume process.

Process Number	Process Description	Status
4	Audit Volume (Inspect Only)	Storage Pool BACKUPPOOL, Volume ADSM.STORAGE.POOL001, Files Processed: 680, Irretrievable Files Found: 0, Partial Files Skipped: 0

To display the results of a volume audit after it has completed, you can issue the QUERY ACTLOG command.

## Auditing Multiple Volumes in a Sequential Access Storage Pool

When you audit a sequential storage volume containing files that span multiple volumes, the server selects all associated volumes and begins the audit process with the first volume on which the first file resides. For example, Figure 53 shows five volumes defined to ENGBACK2. In this example, FileA spans across VOL1 and VOL2, and FileD spans across VOL2, VOL3, VOL4, and VOL5.




Figure 53. Auditing Tape Volumes

If you want to audit volume VOL3, the server first accesses volume VOL2, because FileD begins at VOL2. When volume VOL2 is accessed, the server *only* audits FileD. It does not audit the other files on this volume.

Because FileD spans across multiple volumes, the server accesses volumes VOL2, VOL3, VOL4, and VOL5 to ensure that there are no inconsistencies between the database and the storage pool volumes.

For volumes that require manual mount and dismount operations, this process may entail significant manual intervention.

### **Auditing a Single Volume in a Sequential Access Storage Pool**

To audit a single volume in a sequential storage pool, you can request that the server skip any files that span across the single volume to other volumes in the storage pool. This process is useful when a single volume is damaged, yet you want to audit a different volume that has a file which spans onto the damaged volume.

For example, to audit only volume VOL5 and have the server fix any inconsistencies found between the database and the storage volume, enter:

```
audit volume vol5 fix=yes skippartial=yes
```

## **Requesting Information about Storage Pool Volume Contents**

Any administrator can request information about the contents of a storage pool volume. Viewing the contents of a storage volume is useful when a volume is damaged or before you:

- Request the server to correct any inconsistencies
- Move files from one volume to other volumes
- Delete a volume from a storage pool

Because ADSM tracks the contents of a storage volume through its database, the requested volume need not be accessed in order to determine its contents.

The report generated by a QUERY CONTENT command shows the contents of a volume. This report can be extremely large and may take a long time to produce. To reduce the size of this report, narrow your search by selecting one or all of the following search criteria:

#### **Node name**

Name of the node

#### **File space name**

Remember that file space names are case-sensitive and must be entered by using the exact capitalization used when the file space name was initially defined. Use the QUERY FILESPACE command to find out the correct capitalization.

#### **Number of files to be displayed**

Enter a positive integer, such as 10, to list the first ten files stored on the volume. Enter a negative integer, such as -15, to list the last fifteen files stored on the volume.

**Filetype**

Specifies which types of files.

**Format of how the information is displayed**

Standard or detailed information for the specified volume.

**Damaged**

Specifies whether to restrict the query output either to files that are known to be damaged, or to files that are not known to be damaged.

**Copied**

Specifies whether to restrict the query output to either files that backed up to a copy storage pool, or to files that are not backed up to a copy storage pool.

**Viewing a Standard Report on the Contents of a Volume**

To view the first seven backup files on volume DB0001 from file space /usr on client node TOMC, for example, enter:

```
query content db0001 node=tomc filespace=/usr count=7 type=backup
```

Figure 54 displays a standard report that shows the first seven files from file space /usr on TOMC stored in DB0001.

Node Name	Type	Filespace Name	Client's Name for File
TOMC	Bkup	/usr	/bin/ acctcom
TOMC	Bkup	/usr	/bin/ acledit
TOMC	Bkup	/usr	/bin/ ac1put
TOMC	Bkup	/usr	/bin/ admin
TOMC	Bkup	/usr	/bin/ ar
TOMC	Bkup	/usr	/bin/ arcv
TOMC	Bkup	/usr	/bin/ banner

Figure 54. A Standard Report on the Contents of a Volume

**Viewing a Detailed Report on the Contents of a Volume**

To query the server to display detailed information about the last three files stored on volume VOL1, enter:

```
query content vol1 count=-3 format=detailed
```

Figure 55 on page 225 displays a detailed report that shows the last three files, in reverse order, stored on VOL1. For example, the *test.scr* file is the last file stored on the volume. The segment number, 1/2, identifies that this is the first volume on which *test.scr* resides. The file spans to a second tape volume.

For disk volumes, the *Cached copy?* field identifies whether the file is a cached copy of a migrated file.

```

Node Name: PEASE
Type: Bkup
Filespace Name: /home
Client's Name for File: /pease/dir1/code/ut1/ test.scr
Stored Size: 435
Segment Number: 1/2
Cached Copy?: No

Node Name: PEASE
Type: Bkup
Filespace Name: /home
Client's Name for File: /pease/dir1/code/ut1/ header.scr
Stored Size: 514
Segment Number: 1/1
Cached Copy?: No

Node Name: PEASE
Type: Bkup
Filespace Name: /home
Client's Name for File: /pease/dir1/code/ut1/ appl.scr
Stored Size: 1,013
Segment Number: 1/1
Cached Copy?: No

```

Figure 55. Viewing a Detailed Report of the Contents of a Volume

## Moving Files from One Volume to Another Volume

You can move files from one volume to another volume in the same or a different storage pool. The volumes can be onsite volumes or offsite volumes.

Task	Required Privilege Class
Move files from a volume in any storage pool to an available volume in any storage pool	System or unrestricted storage
Move files from one volume to an available volumes in any storage pools to which you are authorized	Restricted storage

**Note:** Files in a copy storage pool do not move when primary files are moved.

## Moving Data to Other Volumes in the Same Storage Pool

Moving files from one volume to other volumes in the same storage pool is useful:

- When you want to free up all space on a volume so that it can be deleted from the ADSM server

See “Deleting Data Storage Pool Volumes” on page 228 for information about deleting backed up, archived, or client-migrated data before you delete a volume from a storage pool.

- To salvage readable files from a volume that has been damaged

If the server detects a volume write error, it automatically changes the access mode to read-only.

- When you want to delete cached files from disk volumes

If you want to force the removal of cached files, you can delete them by moving data from one volume to another volume. During the move process, ADSM deletes cached files remaining on disk volumes.

If you move data between volumes within the same storage pool and you run out of space in the storage pool before all data is moved from the target volume, then you cannot move all the data from the target volume. In this case, consider moving data to available space in another storage pool as described in “Moving Data to Another Storage Pool.”

## Moving Data to Another Storage Pool

A user might want to move all data from a volume in one storage pool to volumes in another storage pool. When you specify a target storage pool that is different than the source storage pool, ADSM uses the storage hierarchy to move data if more space is required.

**Note:** Data cannot be moved from a primary storage pool to a copy storage pool. Data in a copy storage pool cannot be moved to any other storage pool.

You can move data from random access storage pools to sequential access storage pools. For example, if you have a damaged disk volume and you have a limited amount of disk storage space, you could move all files from the disk volume to a tape storage pool. Moving files from a disk volume to a sequential storage pool may require many volume mount operations. Ensure that you have sufficient personnel and media to move files from disk to sequential storage.

For example, to move the files stored in volume ADSM.STORAGE.POOL001 to any available volume in the BACKTAPE storage pool, enter:

```
move data adsm.storage.pool001 stgpool=backtape
```

When you move data from a volume, the server starts a background process and sends informational messages, such as:

```
ANR1140I Move Data process started for volume ADSM.STORAGE.POOL001  
(process ID 5).
```

During the data movement process, the server:

- Moves any readable files to available volumes in the specified destination storage pool
- Deletes any cached copies from a disk volume
- Attempts to bypass any files that it cannot read

During the data movement process, users cannot access the volume to restore or retrieve files, and no new files can be written to the volume.

## Moving Data from an Offsite Volume in a Copy Storage Pool

The MOVE DATA command can be used to move files from an offsite volume in a copy storage pool.

When you use this command to move files from a volume marked as offsite, ADSM:

- Determines which files are still active on the volume to be reclaimed
- Obtains these files from a primary storage pool or from another copy storage pool
- Copies the files to one or more volumes in the destination copy storage pool.

## Preparing to Move Data

Before you move files from one volume to another volume, complete the following steps:

1. Change the access mode of the storage pool to which the volume belongs to read-only. This process prevents users from backing up or archiving files to volumes within the storage pool.

See “Defining or Updating Storage Pools” on page 194 for information about updating the access mode of a storage pool.

2. Ensure sufficient available space is available on volumes within the specified destination storage pool by:

- Querying the source storage volume to determine how much space is required on other volumes.

See “Monitoring the Use of Storage Pool Volumes” on page 217 for information about requesting information about a storage volume.

- Querying the specified destination storage pool to ensure there is sufficient capacity to store the files being moved. See “Monitoring the Use of Storage Pool Space” on page 202 for information about querying a storage pool.

3. If you need more storage space, define volumes or increase the maximum number of scratch volumes in the specified destination storage pool.

See Chapter 11, “Managing Storage Pool Volumes” on page 213 for preparing volumes to be used for data storage.

4. If you move files from a volume in a sequential storage pool to another volume in the same storage pool, ensure that the mount limit of the device class associated with the storage pool is greater than one.

See “Requesting Information about a Device Classes” on page 168 for requesting information about the mount limit value for the device class.

5. Because moving data from tape to tape requires two tape drives, ensure the tape drives are available.

## Requesting Information about the Data Movement Process

To request information on the data movement process, enter:

```
query process
```

The following figure displays an example of the report that you receive about the data movement process.

Process Number	Process Description	Status
5	Move Data	Storage Pool BACKUPPOOL, Volume ADSM.STORAGE.POOL001, Target Pool BACKTAPE Moved Files: 0, Moved Bytes: 0, Unreadable Files: 0, Unreadable Bytes: 0

## Monitoring the Movement of Data between Volumes

You can query the server for volume information to monitor the movement of data between volumes. For example, to see how much data has moved from the source volume in a move operation (ADSM.STORAGE.POOL001) to available volumes in the BACKTAPE storage pool, enter:

```
query volume * stgpool=back*
```

Figure 56 shows all defined volumes in storage pools beginning with the name BACK. This example shows that the BACKTAPE storage pool is filling as files are moved from ADSM.STORAGE.POOL001 to available volumes in the backup tape storage pool.

Volume Name	Storage Pool Name	Device Class	Estimated Capacity (MB)	%Util	Volume Status
ADSM.STORAGE.POOL1	BACKUPPOOL	DISK	40.0	36.2	On-Line
DB0001	BACKTAPE	CARTRIDGE	180.0	14.3	Filling
DB0002	BACKTAPE	CARTRIDGE	0.0	0.0	Empty

Figure 56. Volume Information Showing Data Movement

## Deleting Data Storage Pool Volumes

Task	Required Privilege Class
Delete volumes from any storage pool	System or unrestricted storage
Delete volumes from storage pools over which they have authority	Restricted storage



Files in a copy storage pool are never deleted unless:

- The volume that contains the copy file is deleted by using the DISCARDDATA=YES option.
- A data-integrity error is detected by using AUDIT VOLUME with the FIX=YES option for a copy storage pool volume.
- The primary file is deleted because:
  - Policy-based file expiration
  - Filespace deletion
  - Deletion of the primary storage pool volume

In addition, if non-cached files are deleted from a primary storage pool volume, any copies of these files in copy storage pools will also be deleted.

## Deleting a Storage Pool Volume with Data

To prevent you from accidentally deleting backed up, archived, or client-migrated files from data storage, the server does not allow you to delete a volume without explicitly identifying whether or not you want to discard user data. The only exception to this rule is if the administrator specifies the DISCARDDATA=YES option on the DELETE VOLUME command.

Before you can delete a storage volume, you must do one of the following:

- Move files to another volume
- Explicitly request that the server discard files from the storage volume

See “Moving Files from One Volume to Another Volume” on page 225 for information about moving data from one volume to another volume.

## Deleting an Empty Storage Pool Volume

For example, to delete an empty volume named ADsm03, enter:

```
delete volume adsm03
```

On an administrative client, you will receive the following confirmation messages, unless the client is running with the NOConfirm option:

```
ANR220W This command will delete volume ADsm03
from its storage pool after verifying that the volume
contains no data.
Do you wish to proceed? (Y/N)
```

After you respond yes, the server generates a background process to delete the volume.

## Deleting a Storage Pool Volume That Contains Residual Data

Even after you move data, residual data may reside on the volume because of I/O or integrity errors. To delete any volume that contains residual data that cannot be moved, you must explicitly specify that files should be discarded from the volume.

For example, to discard all data from volume ADSTM03 before deleting the volume from its storage pool, enter:

```
delete volume adsm03 discarddata=yes
```

When a volume is deleted, the server does not have to access it, because the DELETE VOLUME command only updates the server database.

The server generates a background process and deletes data in a series of batch database transactions. If the delete volume process is cancelled or if a system failure occurs, the volume might still contain data. Reissue the DELETE VOLUME command and explicitly request the server to discard the remaining files on the volume.

After all files have been deleted from the volume, the server deletes the volume from the storage pool.

---

## Restoring Storage Pool Volumes

An administrator can recreate files in primary storage pool volumes using copies in a copy storage pool by issuing the RESTORE VOLUME command.

Task	Required Privilege Class
Define or update volumes in any storage pool for which they have authority	System, unrestricted storage, or restricted storage

Use the RESTORE VOLUME command to restore files that:

- Are currently stored on one or more volumes in the same primary storage pool
- Were previously copied to one or more copy storage pools by using the BACKUP STGPOOL command

If more than one volume is specified to be restored, this command attempts to minimize volume mounts for the copy storage pool. Therefore, to restore more than one volume in the same primary storage pool, issue this command once and specify a list of volumes to be restored.

**Note:** Cached copies of files are never restored. Any cached files that reside on a volume that is being restored are removed from the database during restore processing.

This command changes the access mode of the specified volumes to DESTROYED.

After files are restored, the old references to these files in the primary storage pool are deleted from the database. This means that ADSM will now locate these files on the volumes to which they were restored, rather than on the volume on which

they were previously stored. If a destroyed volume becomes empty because all files have been restored to other locations, the destroyed volume is automatically deleted from the database.

If the backup file copies are moved or deleted during restore processing, the restore process may be incomplete. Therefore, do not issue the following commands for copy storage pool volumes while restore processing is in progress:

- MOVE DATA
- DELETE VOLUME (DISCARDDATA=YES)
- AUDIT VOLUME (FIX=YES)

In addition, you can delay reclamation processing for your copy storage pools by setting the RECLAIM parameter to 100 (using the UPDATE STGPOOL command) while restore processing is in progress.

This command generates a background process that can be canceled with the CANCEL PROCESS command. If a RESTORE VOLUME background process is canceled, some files may have already been restored prior to the cancellation. To display information on background processes, use the QUERY PROCESS command.

When using the RESTORE VOLUME command, be prepared to supply the following information:

**Volume name**

Specifies the name of the volume in the primary storage pool that is being restored.

**Copy storage pool name**

Specifies the name of the copy pool from which the files are to be restored.

**New storage pool**

Specifies the name of the new primary storage pool to which to restore the files.

**Maximum number of processes**

Specifies the maximum number of parallel processes that are used for restoring files.

**Preview**

Specifies whether you want to preview the restore operation before it is actually performed.

See “Recovery from Media Loss” on page 284 for an example of using the RESTORE VOLUME command.



---

## Chapter 12. Managing Tape Operations

ADSM provides server options and commands to help administrators with system or operator privilege control how tapes are used. Running in mount mode lets you monitor tape operations (see “Using Mount Mode” on page 23 for details). You can also use a tape management system to control how the server uses tapes.

---

### Understanding the Server Tape Processing

In order to integrate ADSM tapes into your tape library, you need to understand how the server uses tapes. You can either predefine tape volumes to a storage pool or specify that ADSM dynamically acquire scratch tape volumes, as needed.

ADSM automatically reuses predefined tape volumes when all space on a volume becomes available, such as when reclamation occurs. When scratch tapes become empty, ADSM dynamically deletes the volumes from the storage pool and returns the tapes to a scratch pool.

### Processing Tapes

ADSM uses the normal tape mounting procedures of MVS to mount, dismount, and process tapes. If you have a tape management system installed, you can use that system to manage the server tape processing as you would manage any other tape user.

You can also protect ADSM tapes from being used by other applications by setting the expiration date to 99365 for the device class. See “File Expiration Date” on page 165 for information on setting an expiration date.

In addition, consider using RACF to protect predefined volumes from unauthorized users. See “Protection Status” on page 166 for information on setting protection for tape volumes.

The tape management system is signalled when a tape has been deleted from the ADSM server database through a tape deletion exit. See “Setting Up Tape Installation-Wide Exits” on page 237 for information on using the deletion exit.

### Understanding How the Server Selects and Handles Tapes

In order to effectively manage tape operations, you need to understand how the server selects tapes, requests tape mounts, and maintains information on the condition of tapes. See “How ADSM Stores Files in a Storage Pool Hierarchy” on page 179 for information on how the server selects tapes based on your storage pool definitions.

#### Controlling Tape Mount Request Messages

The server sends messages to request a tape mount. You can define how long a tape can sit idle on a drive.

When defining a device class, you can define how long the server allows a mounted tape to sit idle on the tape drive with the MOUNTRETENTION parameter.

Both the MOUNTWAIT and MOUNTRETENTION parameters may be changed when updating the device class. You can change the other tape options by updating the server options file.

## Requesting Tapes

The server requests tapes by using the standard MVS services for mounting tapes. Handle the server mount requests as you would handle any other tape mount request.

## Maintaining Information on Tape Condition

If the server cannot process a tape correctly, it records an error state for the tape so that it knows the tape cannot be processed correctly. General errors are recorded in the activity log; severe errors are recorded in the error log.

The two error states are *I/O error* and *unavailable*.

A volume is placed in *I/O error* state whenever the server receives a permanent I/O error when processing the tape. A tape in I/O error is no longer written to although the tape can be mounted when a file needs to be retrieved from the tape.

A volume is placed in *unavailable* state when the server has detected the volume has become corrupted, when the CANCEL REQUEST PERMANENT command has been issued to cancel a mount request for the volume, or when the volume status is changed to unavailable with the UPDATE VOLUME command. A tape in unavailable error state is no longer requested by the server, either for writing or reading.

You can use the QUERY VOLUME command to show the error states of volumes. To try to recover files on unavailable volumes, use the following procedure:

1. Use the QUERY CONTENT command to determine what files should be on the tape. If the files are out-of-date or otherwise not needed, proceed to step 6.
2. Use the UPDATE VOLUME command to make the tape available for mounting by setting ACCESS to read-only.
3. Try to read the tape. If the tape can be read without error, it might have been inadvertently marked as unavailable, as a result of a CANCEL REQUEST PERMANENT command. Ensure the volume status is changed to read-only or read-write, and do not proceed with the following steps.
4. Use the MOVE DATA command to try to recover the files that you still need.
5. If the MOVE DATA command does not work, use the AUDIT VOLUME command to synchronize the database with the readable contents of the tape. Unreadable files are deleted from the database, and you cannot recover them from the tape.
6. Use the DELETE VOLUME command to delete the volume from the server's database. If important files are lost, inform the affected users that their files are no longer available.

---

## Handling Tape Mount Requests

You can use the QUERY MOUNT command to determine which tapes are mounted for the server's use, and DISMOUNT VOLUME or DISMOUNT DEVICE commands to dismount them. Use your local procedures to handle any other tape management tasks.

## Querying Mounted Tapes

You can request information about the status of tapes that are mounted with the QUERY MOUNT command. For example, if you enter the following command, the volume name and device address of all mounted tapes is returned to you:

```
query mount
```

This command also returns information on whether the tape is in use or is idle.

If you want information about a mounted tape and you know the volume name of the tape, you can include the name on the command. For example, to get information on mounted tape dsm988, enter:

```
query mount dsm988
```

## Dismounting Tapes

You can dismount tapes by using the DISMOUNT VOLUME command. For example, to dismount the tape volume named VOL001, enter:

```
dismount volume vol001
```

You can use the QUERY MOUNT command to determine the volume names of mounted volumes.

If you know the real address of the tape device on which the volume is mounted (for example, FE8), you can dismount the tape by dismounting the device:

```
dismount device fe8
```

---

## Setting Up ADSM with Your Tape Management System

If you are using a tape management system, you can set up ADSM to use scratch tapes from the scratch pool controlled by the tape management system. Although tape management systems vary, the procedure for setting up ADSM to work with them follows the same basic steps:

- Define ADSM to the tape management system as an external data manager, if applicable.

- Specify to the tape management system that an ADSM volume must be permanently retained.
- Set up tape installation-wide exits

When ADSM requires a new scratch volume, it asks for a private mount. The tape management system can then give ADSM a volume from the scratch pool. However, in most installations, the operator selects the scratch volume from a list provided by the tape management system. The tape management system verifies that the volume mounted is a scratch volume. ADSM uses the volume serial number of this tape to define the volume to the storage pool. The tape label is rewritten with the expiration date from the device class, if it was specified, and the data set name.

When the tape volume becomes empty, usually due to reclamation, it is deleted from ADSM. During the deletion process, the tape deletion exit is called and notifies the tape management system that it can return the tape to the scratch pool. See “Setting Up Tape Installation-Wide Exits” on page 237 for details about the exit.

## Defining ADSM to Your Tape Management System

If you are using IBM's DFSMSrmm or DFRMM tape management system products, you do not have to define ADSM to either of them. The Removable Media Manager is designed to support all tape data sets, including those created by ADSM and DFSMSHsm.

If you are using a different tape management system product, refer to the product documentation. Defining ADSM to tape management systems varies by product, and it is important that you understand how to complete this task for the one being used by your installation.

If your tape management system uses program names to identify External Data Managers, the ADSM program name is ANRSERV.

## Specifying Permanent Tape Retention

You might need to tell the tape management system that a volume must be permanently retained.

If you are using IBM's DFSMSrmm or DFRMM tape management system products, you should define policies (vital record specification's) to cover the ADSM tape data sets. A single policy specifying the ADSM high level qualifier is sufficient to include all ADSM tape data sets.



You can retain all DFSMShsm tapes that require no movement, with the exception of tapes written by ABARS, by specifying the RMM TSO subcommands shown in following example:

```
RMM ADDVRS DSNAME('mprefix.**') COUNT(99999) OWNER(owner)
RMM ADDVRS DSNAME('bprefix.**') COUNT(99999) OWNER(owner)
RMM ADDVRS DSNAME('authid.**') COUNT(99999) OWNER(owner)
```

where:

**mprefix** Specifies the DFSMShsm-defined migrated data set prefix.  
**bprefix** Specifies the DFSMShsm-defined backup and dump data set prefix.  
**authid** Specifies the DFSMShsm prefix used for control data set backups.  
**COUNT(99999)** Specifies to retain all data sets that match the filter mask.  
**owner** Specifies the user ID of the individual responsible for these definitions.

For more information on defining policies, refer to *DFSMS/MVS DFSMSrmm Implementation and Customization Guide* (SC26-4932). For more information on vital record specifications, refer to *DFSMSrmm Guide and Reference* (SC26-4931).

If you are using a different tape management system, one of the most commonly used ways to specify permanent tape retention is to set up the expiration date in the ADSM device class definition.

You should keep the following in mind when setting up the expiration date:

- The storage group to which your tape volumes belong must have this device class.
- An expiration date of 99365 is most often used to indicate permanent retention.
- You can use any value that specifies permanent retention to the tape management system.

You can update a cartridge device class for permanent retention by entering:

```
update devclass cartridge expiration=99365
```

---

Product-sensitive programming interface

## Setting Up Tape Installation-Wide Exits

You must use the deletion installation-wide exit to notify your tape management system that the ADSTAR Distributed Storage Manager server has deleted a tape from its database. This exit can provide a tape management system with the information required to maintain an accurate inventory of ADSM tapes. However, you do not have to use this exit to use tapes with ADSM.

The deletion exit is called when the ADSM server is returning a tape to scratch after it no longer contains valid files, and when the DELETE VOLUME command is used to delete a volume from the database. The tape volume is deleted from the ADSM server database before the deletion exit is called.

### Identifying the Deletion Exit to the Server

In order to use the deletion exit, include the DELETIONEXIT option in the server options file. Specify the name of the module on the option. If you are using a tape management system, specify the module name of the deletion exit provided by the tape management system. For example, if you are using the DFSMSHsm ARCTVEXT module, specify:

```
DELetionexit ARCTVEXT
```

### Processing Requirements for the Deletion Exit

The deletion exit allows your tape management system to accurately reflect the server tape usage. You need to use the tape management system facilities to notify itself that the server has deleted the tape from its database. After the deletion exit sends the notification to the tape management system, it returns control to the server. Tape deletions cannot be cancelled.

If you are using DFSMSrmm or DFRMM, the vital record specification policy is used to determine the management of the volume. If you are using a different tape management system, it should recognize that the volume is being managed by an external data manager (EDM). It should then insure that only ADSM has access until the volume is returned to the tape management system's control via the DELETIONEXIT.

### Writing the Deletion Exit

If you are using the DFSMSHsm ARCTVEXT exit or a deletion exit supplied with a tape management system, include the coding required for the ADSM deletion exit in your existing exit. The coding considerations for this exit are similar to those for the DFSMSHsm ARCTVEXT exit, including entry and exit linkage. A non-zero return code or an abend in the exit results in an error message but has no other effect on ADSM processing.

The deletion exit must be reentrant and is run in the 24-bit addressing mode (AMODE). You should always return to the caller in the caller's AMODE.

### Registers on Entry to the Deletion Exit

The registers contain the following information on entry to the deletion exit:

Register	Description
1	The address of the parameter list passed to the exit.
13	The address of a 72-byte standard MVS save area.
14	The address to which the exit must return control.
15	The address of the deletion exit.
0, 2–12	These registers contain nothing of use to the exit.

## Deletion Exit Parameter List

Table 18 shows the parameter list passed to the tape deletion exit. Register 1 contains the address of this parameter list.

Table 18. Deletion Exit Parameter List (MVS)

Offset	Length (bytes)	Data Type	Description
0 (X'00')	4	Address of the data area	Pointer to an 8-byte data area containing information about the volume the server has deleted from its database. Table 19 shows the format of the data area.
4 (X'04')	4	Address of the return code	The address of a full word containing the return code that the exit must set after completing processing. The return code is binary, and should be 0 if processing is successful, or any other number if processing is unsuccessful. ADSM generates an error message if the processing is unsuccessful. The high order bit of the return code address is set to 1.

Table 19 shows the format of the data area pointed to by the deletion exit parameter list.

Table 19. Deletion Exit Data Area (MVS)

Offset	Length	Data Type or Bit Pattern	Description
0 (X'00')	6	Character	The volume serial number of the tape volume deleted from the database.
6 (X'06')	2	Binary	Informational flags.
		1... ..	The tape volume is purged from the server's database.
		.1.. ..	The server has protected the files on the volume with expiration dates. This information is determined from the device class associated with the volume.  <b>Note:</b> If the expiration or retention information on the device class has changed after the volume was first used, this informational flag may be incorrect.
		.... .. 1... ..	The DFSMSshm SCRATCHTAPE option was used on the SETSYS TAPEDELETION command for this tape category.
.... .. ..1.	Backup tape.		

## Registers on Exit from the Deletion Exit

When you return control to the server, restore the registers to the values they had on entry.

\_\_\_\_\_ End of Product-sensitive programming interface \_\_\_\_\_



---

## Chapter 13. Exporting and Importing Data

Task	Required Privilege Class
Export: Copy some or all server information to sequential volumes Import: Copy server information from sequential volumes to a server	System
Display information about export and import processes	Any administrator

ADSM provides an export-import facility that allows you to copy all or part of a server to removable media so that data can be transferred to another server.

This section takes you through the task of exporting data to sequential media and importing data to create a new ADSM server. The sections listed in the following table begin at the indicated pages.

Section	Page
<b>Concepts:</b>	
Data that can be exported and imported	241
<b>Tasks:</b>	
Preparing to export or import data	242
Monitoring export and import processes	244
Exporting data to sequential media volumes	247
Importing data from sequential media volumes	251

Most tasks presented in this chapter can be performed by using either the graphical user interface or the command-line interface. Table 11 on page 45 shows whether a task can be performed on the graphical user interface, the command-line interface, or both.

For information about issuing ADSM commands, refer to *ADSM Administrator's Reference*. For help performing a task on the graphical user interface, see the procedure described in "Using Online Help" on page 15.

---

### Data That Can Be Export and Imported

Administrators can export or import the following types of ADSM data:

- Server control information, which includes:
  - Administrator definitions
  - Client node definitions
  - Policy and scheduling definitions
- File data, which consists of file space definitions and authorization rules, and any of the following groups of files from data storage:
  - Backed up files and archive copies
  - Active and inactive versions of backed up files
  - Active versions of backed up files

- Archive copies and active versions of backed up files

Your decision on what information to export depends on why you are exporting that information:

- To copy information to a second server to improve client/server performance, use the EXPORT NODE, EXPORT POLICY, and EXPORT ADMIN commands. For example, when many client nodes access the same server, users contend for communication paths, server resources, and tape mounts during a restore or retrieve operation.

To balance client/server performance, you may want to take one or all of the following actions:

- Move a group of client nodes to a second server
- Move policy definitions associated with these client nodes
- Move administrator definitions for administrators who manage these client nodes

Upon successful completion of an import operation, you can delete file spaces, client nodes, policy objects, scheduling objects and administrators from the source server to reduce contention for server resources.

- To copy data for the purpose of installing a new server, use the EXPORT SERVER command to copy all data to tape volumes.

---

## Preparing to Export or Import Data

Before you export or import data, complete the following tasks:

- Use the PREVIEW parameter to verify what data will be moved
- Prepare sequential media for exporting and importing data

## Using Preview before Exporting or Importing Data

ADSM provides the PREVIEW option on the EXPORT and IMPORT commands. With PREVIEW=YES, the report shows how much data will be transferred without any data being moved. With PREVIEW=NO, the export or import operation is performed.

Issue each EXPORT or IMPORT command with PREVIEW=YES to determine which objects and how much data will be moved. ADSM sends the following types of messages to the server console and activity log for each operation:

**Export** Reports the types of objects, number of objects, and number of bytes that would be copied to sequential media volumes. Use this information to determine how many sequential media volumes you need to prepare for an export operation.

**Import** Reports the number and types of objects found on the sequential media volumes that meet your import specifications, and reports information about any problems that it detects, such as corrupted data. Use this information to determine which data to move from sequential media volumes to the server and to determine if you have enough storage pool space allocated on the server for the import operation to succeed.

To determine how much space is required to export server definitions and all backup versions and archive copies from data storage to sequential media volumes, enter:

```
export server filedata=all preview=yes
```

After you issue this command, the server starts a background process and issues a message similar to the following:

```
EXPORT SERVER started as Process 4
```

You can request information about the background process, as described in “Requesting Information about an Export or Import Process” on page 244. If necessary, you can cancel an export or import process, as described in “Canceling Server Processes” on page 54.

## Planning for Sequential Media Used to Export Data

To export data, you must specify a device class that supports removable media and identify the volumes that will be used to store the exported data. Use this section to help you select the device classes and prepare sequential media volumes.

### Selecting a Device Class

You can query the server about device classes in order to select a device class supported by both the source and target servers. If the existing device classes are not supported on both the source and target server, then define a new device class, as described in Chapter 9, “Managing Storage Devices” on page 161.

**Note:** If the mount limit for the device class selected is reached when you request an export (that is, if all the drives are busy), ADSM automatically cancels lower priority operations, such as reclamation, to make a mount point available for the export.

### Estimating the Number of Tapes to Label

To estimate the number of tapes required to store export data, divide the number of bytes to be moved by the estimated capacity of a volume.

For example, cartridge system tape volumes used with 3490 tape devices have an estimated capacity of 360MB. If the preview shows that you need to transfer 720MB of data, then label at least two tape volumes before you export the data.

### Using Scratch Media

ADSM allows you to use scratch media to ensure that you have sufficient space on which to store all export data. If you use scratch media, be sure to record their label names and the order in which they were mounted.

## Labelling Tapes

During an import process, you must specify the order in which tape volumes will be mounted. This order must match the order in which tapes have been mounted during the export process.

To ensure that tapes are mounted in the correct order, label tapes with information that identifies the order in which they are mounted during the import process. For example, label tapes as DSM001, DSM002, DSM003, and so on to indicate the order in which data is stored on the tape volumes.

When you export data, record the date and time for each labeled tape. Store this information in a safe location, because as you will need the information when you import the data to the server.

---

## Monitoring Export and Import Processes

ADSM provides you with a number of methods for monitoring export or import processes. From the server console or from an administrative client running in console mode, you can display information about:

- A process while it is running
- The activity log for status information when a process has completed

## Requesting Information about an Export or Import Process

After you issue an EXPORT or IMPORT command, the server starts a background process, assigns a process ID to the operation, and displays the process ID when the operation starts.

You can query an export or import process by specifying the process ID number. For example, to request information about the EXPORT SERVER operation, which started as process 4, enter:

```
query process 4
```

If you issue a preview version of an EXPORT or IMPORT command and then query the process, ADSM reports the types of objects to be copied, the number of objects to be copied, and the number of bytes to be copied.

When you export or import data and then query the process, ADSM displays the number and types of objects copied so far, and the total number of bytes that have been transferred, along with information on any media mount requests that may be outstanding for the process.

For guidance information on querying background processes, see “Requesting Information about Server Processes” on page 54.



## Viewing Information from the Server Console

When you issue an IMPORT or EXPORT command, either from the server console or from an administrative client, information is displayed on the server console.

Figure 57 shows an example of the information that is displayed after issuing an EXPORT SERVER command.

```
ANR0610I EXPORT SERVER started by SERVER_CONSOLE as process 1.
ANR0639I EXPORT SERVER: Processing domain ENGPOLDOM.
ANR0640I EXPORT SERVER: Processing policy set ACTIVE in policy domain
ENGPOLDOM.
ANR0640I EXPORT SERVER: Processing policy set STANDARD in policy domain
ENGPOLDOM.
ANR0641I EXPORT SERVER: Processing management class STANDARD in domain
ENGPOLDOM, set ACTIVE.
ANR0641I EXPORT SERVER: Processing management class STANDARD in domain
ENGPOLDOM, set STANDARD.
ANR0643I EXPORT SERVER: Processing archive copy group in domain
ENGPOLDOM, set STANDARD, management class ACTIVE.
ANR0643I EXPORT SERVER: Processing archive copy group in domain
ENGPOLDOM, set STANDARD, management class STANDARD.
ANR0643I EXPORT SERVER: Processing backup copy group in domain
ENGPOLDOM, set STANDARD, management class ACTIVE.
ANR0643I EXPORT SERVER: Processing backup copy group in domain
ENGPOLDOM, set STANDARD, management class STANDARD.
ANR0604I EXPORT SERVER: No schedules were found in policy domain * for
exporting.
ANR0635I EXPORT SERVER: Processing node TOMC.
ANR0605I EXPORT SERVER: No schedule associations were found in
policy domain * for exporting.
ANR0637I EXPORT SERVER: Processing file space DRIVED for node TOMC.
ANR0637I EXPORT SERVER: Processing file space OS2 for node TOMC.
ANR0637I EXPORT SERVER: Processing file space OS2VDISK for node TOMC.
ANR0617I EXPORT SERVER: Processing completed successfully.
ANR0620I EXPORT SERVER: Copied 1 domain(s).
ANR0621I EXPORT SERVER: Copied 2 policy set(s).
ANR0622I EXPORT SERVER: Copied 2 management class(es).
ANR0623I EXPORT SERVER: Copied 4 copy group(s).
ANR0626I EXPORT SERVER: Copied 1 node definition(s).
ANR0627I EXPORT SERVER: Copied 3 file space(s), 16 archive file(s)
and 0 backup file(s).
ANR0629I EXPORT SERVER: Copied 3045632 bytes of data.
ANR0611I EXPORT SERVER started by SERVER_CONSOLE as process 1 has ended.
```

Figure 57. Sample Export Server Output Sent to the Server Console

## Viewing Information from an Administrative Client

Use the console mode from an administrative client to monitor export or import operations or to capture processing messages to an output file.

To start an administrative session in console mode, enter:

```
dsmdmc -consolemode
```

While the system is running in console mode, you cannot enter any administrative commands from the client session. You can, however, start another administrative client session for entering commands (for example, QUERY PROCESS) if you are using a multitasking workstation, like OS/2 or AIX.

If you want ADSM to write all terminal output to a file, specify the OUTFILE option with a destination. For example, to write output to the SAVE.OUT file, enter:

```
dsmdmc -consolemode -outfile=save.out
```

For information about using the CONSOLE mode option and ending an administrative session in console mode, refer to *ADSM Administrator's Reference*.

## Querying the Activity Log for Export or Import Information

After an export or import process has completed, you can query the activity log for status information and possible error messages.

To minimize processing time when querying the activity log for export or import information, restrict the search by specifying *export* or *import* in the SEARCH parameter of the QUERY ACTLOG command.

For example, to determine how much data will be moved after issuing the preview version of the EXPORT SERVER command, query the activity log by entering:

```
query actlog search=export
```

Figure 58 on page 247 displays a sample activity log report.

Date/Time	Message
05/03/1994 10:50:28	ANR0610I EXPORT SERVER started by SERVER_CONSOLE as process 1.
05/03/1994 10:50:28	ANR0639I EXPORT SERVER: Processing domain ENGPOLDOM.
05/03/1994 10:50:28	ANR0640I EXPORT SERVER: Processing policy set ACTIVE in policy domain ENGPOLDOM.
05/03/1994 10:50:28	ANR0640I EXPORT SERVER: Processing policy set STANDARD in policy domain ENGPOLDOM.
05/03/1994 10:50:29	ANR0641I EXPORT SERVER: Processing management class STANDARD in domain ENGPOLDOM, set ACTIVE.
05/03/1994 10:50:29	ANR0641I EXPORT SERVER: Processing management class STANDARD in domain ENGPOLDOM, set STANDARD.
05/03/1994 10:50:29	ANR0643I EXPORT SERVER: Processing archive copy group in domain ENGPOLDOM, set STANDARD, management class ACTIVE.
05/03/1994 10:50:29	ANR0643I EXPORT SERVER: Processing archive copy group in domain ENGPOLDOM, set STANDARD, management class STANDARD.
05/03/1994 10:50:29	ANR0642I EXPORT SERVER: Processing backup copy group in domain ENGPOLDOM, set STANDARD, management class ACTIVE.
05/03/1994 10:50:29	ANR0642I EXPORT SERVER: Processing backup copy group in domain ENGPOLDOM, set STANDARD, management class STANDARD.
05/03/1994 10:50:29	ANR0604I EXPORT SERVER: No schedules were found in policy domain * for exporting.
05/03/1994 10:50:29	ANR0635I EXPORT SERVER: Processing node TOMC.
05/03/1994 10:50:29	ANR0605I EXPORT SERVER: No schedule associations were found in policy domain * for exporting.
05/03/1994 10:50:29	ANR0637I EXPORT SERVER: Processing file space DRIVED for node TOMC.
05/03/1994 10:50:29	ANR0637I EXPORT SERVER: Processing file space OS2 for node TOMC.
05/03/1994 10:50:29	ANR0637I EXPORT SERVER: Processing file space OS2VDISK for node TOMC.
05/03/1994 10:50:32	ANR0617I EXPORT SERVER: Processing completed successfully.
05/03/1994 10:50:32	ANR0620I EXPORT SERVER: Copied 1 domain(s).
05/03/1994 10:50:32	ANR0621I EXPORT SERVER: Copied 2 policy set(s).
05/03/1994 10:50:32	ANR0622I EXPORT SERVER: Copied 2 management class(es).
05/03/1994 10:50:32	ANR0623I EXPORT SERVER: Copied 4 copy group(s).
05/03/1994 10:50:32	ANR0626I EXPORT SERVER: Copied 1 node definition(s).
05/03/1994 10:50:32	ANR0627I EXPORT SERVER: Copied 3 file space(s), 16 export file(s) and 0 backup file(s).
05/03/1994 10:50:32	ANR0629I EXPORT SERVER: Copied 3045632 bytes of data.
05/03/1994 10:50:32	ANR0611I EXPORT SERVER started by SERVER_CONSOLE as process 1 has ended.

Figure 58. Sample Activity Log Report on Exported Data

## Exporting Data to Sequential Media Volumes

You can export all server control information or a subset of server control information by specifying one or more of the following export commands:

- EXPORT SERVER
- EXPORT ADMIN
- EXPORT NODE
- EXPORT POLICY

When you export data, you must specify the device class to which export data can be written. You must also list the volumes in the order in which they are mounted when the data is imported. See "Labelling Tapes" on page 244 for information on labelling tape volumes.

## Deciding When to Export Data to Tape

When you issue an EXPORT command, the operation runs as a background process. This process allows you to continue performing administrative tasks. In addition, users can continue to back up, archive, restore, or retrieve files from ADSM.

If you choose to perform an export operation during normal working hours, be aware that administrators can change server definitions and users may modify backed up or archived files. If administrators or users modify data shortly after it has been exported, then the information copied to tape may not be consistent with data stored on the source server.

If you want to export an exact point-in-time copy of server control information, you can prevent administrative and clients nodes from accessing the server as described in “Preventing Administrative Clients from Accessing the Server.”

### Preventing Administrative Clients from Accessing the Server

Administrators can change administrator, policy, or client node definitions during an export process. To prevent administrators from modifying these definitions, you can lock out administrator access to the server and cancel any administrative sessions before issuing an EXPORT command. After the export process is complete, unlock administrator access.

For more information on canceling sessions, see “Canceling a Client Session” on page 52. For more information on locking or unlocking administrators from the server, see “Locking and Unlocking Administrators from the Server” on page 101.

### Preventing Client Nodes from Accessing the Server

If client node information is exported while the same client is performing a backup or archive, the latest file copies for the client may not be exported to tape. To prevent users from accessing the server during export operations, cancel existing client sessions as described in “Canceling a Client Session” on page 52. Then you can do one of the following:

- Disable server access to prevent client nodes from accessing the server, as described in “Disabling or Enabling Server Access” on page 53.

This option is useful when you export all client node information from the source server and want to prevent all client nodes from accessing the server.

- Lock out particular client nodes from server access, as described in “Locking and Unlocking Client Nodes” on page 106.

This option is useful when you export a subset of client node information from the source server and want to prevent particular client nodes from accessing the server until the export operation is complete.

After the export operation is complete, allow client nodes to access the server again by:

- Enabling the server, as described in “Disabling or Enabling Server Access” on page 53.
- Unlocking client nodes, as described in “Locking and Unlocking Client Nodes” on page 106

## Exporting Server Data

When you issue the EXPORT SERVER command, the server exports all server control information. You can also export file data information with the EXPORT SERVER command.

For example, you want to export definitions and all file data to four defined tape cartridges, which are supported by the CARTRIDGE device class. To ensure that you have sufficient space on tape volumes, allow ADSM to use scratch volumes by using the default of SCRATCH=YES. To issue this command, enter:

```
export server devclass=cartridge -  
volumenames=dsm001,dsm002,dsm003,dsm004 filedata=all
```

During the export process, ADSM exports definition information before it exports file data information. This ensures that definition information is stored on the first tape volumes. This process allows you to mount a minimum number of tapes during the import process, if your goal is to copy only control information to the target server.

In the example above, the server exports:

- Administrator definitions
- Client node definitions
- Policy domain, policy set, management class, and copy group definitions
- Schedule definitions and client node associations
- File space definitions
- File space authorization rules
- Backed up, archived, and space-managed files

## Exporting Administrator Information

When you issue the EXPORT ADMIN command, the server exports administrator definitions. Each administrator definition includes:

- Administrator name, password and contact information
- Any administrative privilege classes the administrator has been granted
- Whether the administrator ID is locked from server access

You can specify a list of administrator names, or you can export all administrator names from the server to tape volumes.

In the following example, definitions for the DAVEHIL and PENNER administrator IDs will be exported to the DSM001 tape volume, which is supported by the CARTRIDGE device class. Do not allow any scratch media to be used during this export process. To issue this command, enter:

```
export admin davehil,penner devclass=cartridge -  
volumenames=dsm001 scratch=no
```

## Exporting Client Node Information

When you issue the EXPORT NODE command, the server exports client node definitions. Each client node definition includes:

- User ID, password, and contact information
- Name of the policy domain to which the client is assigned
- File compression status
- Whether the user has the authority to delete backed up or archived files from data storage
- Whether the client node ID is locked from server access

Optionally, you can specify whether to export file data. File data consists of file space definitions and authorizations, and any of the following groups of files:

- Backed up files and archive copies
- Active and inactive versions of backed up files
- Active versions of backed up files
- Archive copies
- Archive copies and active versions of backed up files

When client file data is exported, ADSM copies files to export volumes in the order of their physical location in data storage. This process minimizes the number of mounts required during the export process.

If you do not explicitly specify that you want to export file data, then ADSM only exports client node definitions.

In the following example, definitions for client nodes and file spaces in the ENGPOLDOM policy domain will be exported. Also, export any active backup versions of files belonging to these client nodes. Finally, export this information to scratch volumes, which are supported by the CARTRIDGE device class. To issue this command, enter:

```
export node filespace=* domains=engpoldom -  
filedata=backupactive devclass=cartridge
```

In this example, the server exports:

- Definitions of client nodes assigned to the engineering policy domain
- File space definitions and backup authorizations for each client node in the engineering policy domain
- Active versions of backed up files belonging to the client nodes assigned to the engineering policy domain

## Exporting Policy Information

When you issue the EXPORT POLICY command, the server exports the following information belonging to each specified policy domain:

- Policy domain definitions
- Policy set definitions, including the active policy set
- Management class definitions, including the default management class
- Backup copy group and archive copy group definitions
- Schedule definitions

- Associations between client nodes and schedules

In the following example, policy and scheduling definitions in the policy domain named ENGPOLDOM will be exported. Identify tape volumes DSM001 and DSM002, which belong to the CARTRIDGE device class. Finally, allow the server to use scratch tape volumes if necessary by using the default of SCRATCH=YES. To issue this command, enter:

```
export policy engpoldom -  
devclass=cartridge volumenames=dsm001,dsm002
```

---

## Importing Data from Sequential Media Volumes

Before you import data to a new target server, a system programmer must:

1. Install ADSM from program sequential media on the target server.

Refer to *ADSM Installing the Server and Administrative Client* for information on installing ADSM.

2. Define disk space for the database and recovery log.

For information on defining space, refer to *ADSM Installing the Server and Administrative Client*.

3. Define storage pools and volumes.

Because each server operating system requires different naming conventions for volumes used by storage pools, ADSM does not export data storage definitions. Therefore, you must define initial storage pools and volumes on the target server, as described in the *ADSM Installing the Server and Administrative Client*.

After ADSM is installed and set up on the target server, a system administrator can import all server control information or a subset of server control information by specifying one or more of the following import commands:

- IMPORT SERVER
- IMPORT ADMIN
- IMPORT NODE
- IMPORT POLICY

This section guides you through the entire process of importing all server control information and file data from tape volumes to a new target server. This process includes:

- Previewing information before you import data
- Importing definitions
- Tailoring data storage definitions on the target server
- Importing file data

After you understand how to import server control information and file data information, you can import any subset of data to the target server.

## Step 1: Previewing Information before You Import Data

Before you import any data to the target server, preview each import command to determine what data you want to import to the target server. You can import all or a subset of export data from tapes.

When you set PREVIEW=YES, tape operators must mount export tape volumes so that the target server can calculate the statistics reported by the use of this parameter.

For example, to preview information for the IMPORT SERVER command, enter:

```
import server devclass=cartridge preview=yes -  
volumenames=dsm001,dsm002,dsm003,dsm004
```

Figure 59 on page 253 shows an example of the messages sent to the server console and activity log.



```

ANR0402I Session 3 started for administrator SERVER_CONSOLE (Server).
ANR1363I Import volume DSM001 opened (sequence number 1).
ANR0610I IMPORT SERVER started by SERVER_CONSOLE as process 2.
ANR0612I IMPORT SERVER: Reading EXPORT SERVER data from server ADSM exported
05/07/1994 12:39:48.
ANR0639I IMPORT SERVER: Processing domain ENGPOLDOM.
ANR0640I IMPORT SERVER: Processing policy set ACTIVE in policy domain
ENGPOLDOM.
ANR0640I IMPORT SERVER: Processing policy set STANDARD in policy domain
ENGPOLDOM.
ANR0641I IMPORT SERVER: Processing management class STANDARD in domain
ENGPOLDOM, set ACTIVE.
ANR0641I IMPORT SERVER: Processing management class MCENG in domain
ENGPOLDOM, set STANDARD.
ANR0641I IMPORT SERVER: Processing management class STANDARD in domain
ENGPOLDOM, set STANDARD.
ANR0643I IMPORT SERVER: Processing archive copy group in domain
ENGPOLDOM, set ACTIVE, management class STANDARD.
ANR0643I IMPORT SERVER: Processing archive copy group in domain ENGPOLDOM,
set STANDARD, management class MCENG.
ANR0643I IMPORT SERVER: Processing archive copy group in domain ENGPOLDOM,
set STANDARD, management class STANDARD.
ANR0642I IMPORT SERVER: Processing backup copy group in domain ENGPOLDOM,
set ACTIVE, management class STANDARD.
ANR0642I IMPORT SERVER: Processing backup copy group in domain ENGPOLDOM,
set STANDARD, management class MCENG.
ANR0642I IMPORT SERVER: Processing backup copy group in domain ENGPOLDOM,
set STANDARD, management class STANDARD.
ANR0638I IMPORT SERVER: Processing administrator DAVEHIL.
ANR0638I IMPORT SERVER: Processing administrator PENNER.
ANR0635I IMPORT SERVER: Processing node TOMC.
ANR0636I IMPORT SERVER: Processing file space OS2 for node TOMC as file
space OS1.
ANR0636I IMPORT SERVER: Processing file space DRIVED for node TOMC as file
space DRIVE1.
ANR0636I IMPORT SERVER: Processing file space OS2VDISK for node TOMC as file
space OS2VDIS1.
ANR1365I Import volume DSM001 closed (end reached).
ANR1363I Import volume DSM002 opened (sequence number 2).
ANR1365I Import volume DSM002 closed (end reached).
ANR1363I Import volume DSM003 opened (sequence number 3).
ANR1365I Import volume DSM003 closed (end reached).
ANR1363I Import volume DSM004 opened (sequence number 4).
ANR1365I Import volume DSM004 closed (end reached).
ANR0617I IMPORT SERVER: Processing completed successfully.
ANR0620I IMPORT SERVER: Copied 1 domain(s).
ANR0621I IMPORT SERVER: Copied 2 policy set(s).
ANR0622I IMPORT SERVER: Copied 2 management class(es).
ANR0623I IMPORT SERVER: Copied 6 copy group(s).
ANR0625I IMPORT SERVER: Copied 2 administrator(s).
ANR0626I IMPORT SERVER: Copied 1 node definition(s).
ANR0627I IMPORT SERVER: Copied 3 file space(s), 0 archive file(s) and 462
backup file(s).
ANR0629I IMPORT SERVER: Copied 8856358 bytes of data.
ANR0611I IMPORT SERVER started by SERVER_CONSOLE as process 2 has ended.

```

Figure 59. Sample Report Created by Issuing Preview for an Import Server Command

Use the value reported for the total number of bytes copied to estimate if you have sufficient storage pool space on the server to store imported file data.

For example, Figure 59 shows that 8856358 bytes of data will be imported. Ensure that you have at least 8856358 bytes of available space in the backup storage pools defined to the server. You can use the QUERY STGPOOL and

QUERY VOLUME commands to determine how much space is available in the server storage hierarchy.

In addition, the preview report shows that 0 archive files and 462 backup files will be imported. Because backup data is being imported, ensure that you have sufficient space in the backup storage pools used to store this backup data. See “Step 3: Tailoring Data Storage Definitions on the Target Server” on page 256 for information on identifying storage pools on the target server.

For information on specifying the PREVIEW parameter, see “Using Preview before Exporting or Importing Data” on page 242. For information on reviewing the results of a preview operation, see “Monitoring Export and Import Processes” on page 244.

## Step 2: Importing Definitions

Next, you want to import server control information, which includes:

- Administrator definitions
- Client node definitions
- Policy domain, policy set, management class, and copy group definitions
- Schedule definitions and client node associations

However, do not import file data at this time, because some storage pools named in the copy group definitions may not exist yet on the target server.

Before you import server control information, decide whether to replace existing definitions. You should also understand how ADSM handles the activation of policy definitions.

After you understand how data is imported to the target server, you can start an administrative client session in console mode to capture import messages to an output file. Then import the server control information from specified tape volumes.

### Determining Whether to Replace Existing Definitions

By using the REPLACEDEFS option, you can specify whether to replace existing definitions on the target server when ADSM encounters an object with the same name during the import process.

For example, if a definition exists for the ENGPOLDOM policy domain on the target server before you import policy definitions, then you must specify REPLACEDEFS=YES to have ADSM replace the existing definition with the data from the export tape.

Definitions that can be replaced include administrator, client node, policy, or schedule definitions. The default is to not replace existing definitions on the target server.

### Understanding How ADSM Imports Active Policy Sets

When ADSM imports policy definitions, the following objects are imported to the target server:

- Policy domain definitions
- Policy set definitions, including the ACTIVE policy set
- Management class definitions
- Backup copy group definitions

- Archive copy group definitions
- Schedule definitions defined for each policy domain
- Client node associations, if the client node definition exists on the target server

If ADSM encounters a policy set named ACTIVE on the tape volume during the import process, it uses a temporary policy set named \$\$ACTIVE\$\$ to import the active policy set.

After \$\$ACTIVE\$\$ is imported to the target server, ADSM activates this policy set. During the activation process, the server validates the policy set by examining the management class and copy group definitions.

ADSM reports on the following conditions, which result in warning messages during validation:

- The storage destinations specified in the backup copy groups and the archive copy groups do not refer to defined storage pools.
- The default management class does not contain a backup or archive copy group.
- The current ACTIVE policy set contains management class names that are not defined in the policy set to be activated.
- The current ACTIVE policy set contains copy group names that are not defined in the policy set to be activated.

After each \$\$ACTIVE\$\$ policy set has been activated, ADSM deletes that \$\$ACTIVE\$\$ policy set from the target server. To view information about active policy on the target server, you can use the following commands:

- QUERY COPYGROUP
- QUERY MGMTCLASS
- QUERY POLICYSET

Results from issuing the QUERY DOMAIN command show the activated policy set as \$\$ACTIVE\$\$\$. ADSM uses the \$\$ACTIVE\$\$ name to show you that the policy set which is currently activated for this domain is the policy set that was active at the time the export was performed.

### Directing Import Messages to an Output File

The information generated by the validation process can help you define a storage hierarchy that supports the storage destinations currently defined in the import data.

You can direct import messages to an output file to capture any error messages that are detected during the import process by starting an administrative client session in console mode before you invoke this import command.

For example, to direct messages to an output file named IMPSERV.OUT, enter:

```
dsadm -consolemode -outfile=impserv.out
```

## Importing Server Control Information

Now you are ready to import the server control information. Based on the information generated during the preview operation, you know that all definition information has been stored on the first tape volume named DSM001. Specify that this tape volume can be read by a device belonging to the CARTRIDGE device class.

From an administrative client session or from the server console, enter:

```
import server filedata=none -  
devclass=cartridge volumenames=dsm001
```

## Step 3: Tailoring Data Storage Definitions on the Target Server

After you import definition information, use the reports generated by the import process to help you tailor your storage hierarchy on the target server.

To tailor data storage definitions on the target server, complete the following steps:

1. Identify any storage destinations defined in copy groups that do not match defined storage pools by:
  - Reviewing the error messages generated during the validation process, which have been directed to an output file on the administrative client using console mode
  - Querying copy group definitions to compare the storage destinations defined in the copy group with the name of storage pools on the target server

To request detailed reports for all backup copy groups and archive copy groups, enter:

```
query copygroup * active * standard type=backup format=detailed  
query copygroup * active * standard type=archive format=detailed
```

2. Ensure that all copy group definitions in ACTIVE policy sets refer to defined storage pools by doing one of the following:
  - Defining storage pools that match the storage destination names for each copy group, as described in “Defining or Updating Storage Pools” on page 194
  - Changing the names of the storage pools referenced by the copy groups in the ACTIVE policy set; refer them to storage pools defined on the target server by:
    - a. Copying the ACTIVE policy set to another policy set
    - b. Modifying the storage destinations of copy groups in that policy set
    - c. Activating the new policy set

For information on copying policy sets, see “Defining and Updating a Policy Set” on page 127.

Depending upon the amount of client file data that you expect to import, you may want to examine the storage hierarchy to ensure that sufficient storage space is

available, and that secondary storage pools are defined for migration if or when the storage pools specified by copy groups fill up with data.

## Step 4: Importing File Data Information

After you have defined the appropriate storage hierarchy on the target server, you can import file data from the tape volumes. File data consists of all file space definitions and authorization rules and any of the following groups of files:

- Backed up files and archive copies
- Active and inactive versions of backed up files
- Active versions of backed up files
- Archive copies
- Archive copies and active versions of backed up files

Before you import file data information:

- Understand how ADSM handles duplicate file space names
- Decide whether to keep the original creation date for backup versions and archive copies or to import file data using an adjusted date.

Then you can import file data to the target server.

### Understanding How Duplicate File Spaces Are Handled

When ADSM imports file data information, it imports any file spaces belonging to each specified client node. If a file space definition already exists on the target server for the node, ADSM does *not* replace the existing file space name.

If ADSM encounters duplicate file space names when it imports file data information, it creates a new file space name for the imported definition by replacing the final character or characters with a number. A message showing the old and new file space names is written to the server console and to the activity log.

For example, if the C\_DRIVE and D\_DRIVE file space names reside on the target server for node FRED and on the tape volume for FRED, then the server imports the C\_DRIVE file space as C\_DRIV1 file space and the D\_DRIVE file space as D\_DRIV1 file space, both assigned to node FRED.

### Deciding Whether to Use a Relative Date When Importing File Data

When you import file data, you can keep the original creation date for backup versions and archive copies, or you can specify that ADSM use an adjusted date.

Because export-import tape volumes might not be used for some time after an export operation, the original dates defined for backup versions and archive copies may be old enough that files are expired immediately by policy when the data is imported to the target server.

To prevent backup versions and archive copies from being expired immediately, specify DATES=RELATIVE on the IMPORT NODE or IMPORT SERVER commands to adjust for the elapsed time since the files were exported to tape.

For example, assume that the export-import tape contains an archive copy archived five days prior to the export operation. If the tape volume resides on the shelf for

six months before the data is imported to the target server, ADSM resets the archival date to five days prior to the import operation.

If you want to keep the original backup and archive dates set for backup versions and archive copies, then use `DATES=ABSOLUTE`, which is the default. If you use the absolute value, then any files whose retention period has passed will be expired shortly after they are imported to the target server.

### **Issuing an Import Server or Import Node Command**

You can import file data, either by issuing the `IMPORT SERVER` or `IMPORT NODE` command. When you issue either of these commands, you can specify which type of files should be imported for all client nodes specified and found on the export tapes. You can specify any of the following values to import file data:

#### **All**

Specifies that all archive and backup copies for specified client nodes are imported to the target server

#### **None**

Specifies that no files are imported to the target server; only client node definitions are imported

#### **Archive**

Specifies that only archive copies are imported to the target server

#### **Backup**

Specifies that only backup copies, whether active or inactive, are imported to the target server

#### **Backupactive**

Specifies that only active backup copies will be imported to the target server

#### **Allactive**

Specifies that only archive copies and active versions of backed up files are to be imported

#### **Spacemanaged**

Specifies that only files that have been migrated from a user's local file system are to be imported

In the following example, you will import all backup and archive copies to the target server. However, do not replace any existing server control information during this import operation. Specify the four tape volumes that were identified during the preview operation. These tape volumes can be read by any device in the `CAR-TRIDGE` device class. To issue this command, enter:

```
import server filedata=all replacedefs=no -  
devclass=cartridge volumenames=dsm001,dsm002,dsm003,dsm004
```

## Considerations When Importing Data

Although you can issue any import command, ADSM only imports information that can be found on the tape volumes mounted during import. In addition, you can use any import command to copy pertinent information from export-import tapes to the target server, no matter what export command was used to create the export-import tapes.

### Importing a Subset of Information from Tapes

Issue the following operations to import a subset of information to a target server:

- An IMPORT ADMIN command against tapes created with the EXPORT SERVER command
- An IMPORT NODE command against tapes created with the EXPORT SERVER command
- An IMPORT POLICY command against tapes created with the EXPORT SERVER command
- An IMPORT SERVER command against tapes created with the EXPORT NODE command
- An IMPORT SERVER command against tapes created with the EXPORT ADMIN command
- An IMPORT SERVER command against tapes created with the EXPORT POLICY command

### Issuing Useless Import Commands

While ADSM allows you to issue any import command, data cannot be imported to the server if it has not been exported to tape. You cannot issue the following operations to import data:

- The IMPORT POLICY command against tapes created with the EXPORT NODE command
- The IMPORT POLICY command against tapes created with the EXPORT ADMIN command
- The IMPORT ADMIN command against tapes created with the EXPORT POLICY command
- The IMPORT ADMIN command against tapes created with the EXPORT NODE command
- The IMPORT NODE command against tapes created with the EXPORT POLICY command
- The IMPORT NODE command against tapes created with the EXPORT ADMIN command

## Recovering from Errors during the Import Process

During import processing, the server may encounter invalid data due to corruption during storage on tape or in the database prior to the export operation. If invalid data is encountered during an import operation, the server does the following:

- If a new object is being defined, the default value is used
- If the object already exists, the existing parameter is not changed

The server reports on the affected objects to the server console and activity log during import and export operations. You should query these objects when the import process is complete to see if they reflect information that is acceptable to you.

Each time you run the `IMPORT NODE` or `IMPORT SERVER` command with the `FILEDATA` parameter equal to a value other than `NONE`, ADSM creates a new file space and imports data to it. This process ensures that the current import does not overwrite data from a previous import. For information on how ADSM handles duplicate file spaces, see “Understanding How Duplicate File Spaces Are Handled” on page 257.

A file space definition may already exist on the target server for the node. If so, an administrator with system privilege can issue the `DELETE FILESPACE` command to remove file spaces that are corrupted or no longer needed. For more information on the `DELETE FILESPACE` command, refer to the *ADSM Administrator's Reference*.

### **Renaming a File Space**

An imported file space can have the same name as a file space that already exists on a client node. In this case, the server does not overlay the existing file space, and the imported file space is given a new system generated file space name. This new name may match file space names that have not been backed up and are unknown to the server. In this case, you can use the `RENAME FILESPACE` command to rename the imported file space to the naming convention used for the client node.



---

## Chapter 14. Recovering Data

The ADSM database, recovery log, and storage pools are critical to the operation of the server. If the database or recovery log are 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.

To help in the recovery of your data if it is lost, you should use the following strategy:

- Mirror the server database and recovery log
- Back up the server database periodically to removable media and store the media offsite
- Back up the storage pools periodically to removable media and store the media offsite

ADSM also provides standalone database salvage utilities that can be used as a final measure (see “Database Salvage Utilities” on page 284 for details).

The sections listed in the following tables begin at the indicated pages.

### MIRRORING

Section	Page
<b>Concepts:</b>	
Database and recovery log mirroring	262
<b>Tasks:</b>	
Using mirroring	263
Recovering by using mirrored copies of the database	265

### DATABASE BACKUP AND RECOVERY

Section	Page
<b>Concepts:</b>	
Database backup and recovery	265
<b>Tasks:</b>	
Using database backup features	269
Recovering by using backed up copies of the database	276

## STORAGE POOL BACKUP AND RECOVERY

Section	Page
<b>Concepts:</b>	
Storage pool backup and recovery	279
<b>Tasks:</b>	
Using storage pool backup features	280
Recovering by using backed up copies of storage pools	280

## BACKUP AND RECOVERY EXAMPLES

Section	Page
ADSM backup and recovery scenarios	281
Backup	281
Recovery from a disaster	282
Recovery from media loss	284

Most tasks presented in this chapter can be performed using either the graphical user interface or the command line interface. Table 12 on page 46 shows whether a task can be performed on the graphical user interface, the command line interface, or both.

For information about issuing ADSM commands, refer to *ADSM Administrator's Reference*. For help performing a task on the graphical user interface, see the procedure described in "Using Online Help" on page 15.

---

## Database and Recovery Log Mirroring

One of the most basic but catastrophic failures in ADSM is the loss of the database or recovery log due to a hardware failure. ADSM helps to protect against this by mirroring the database and recovery log, which writes the same data to multiple disks simultaneously. However, mirroring does not protect against a hardware failure that affects multiple drives or causes the loss of the entire system. Administrator control of mirroring (starting, stopping, and so on) can be performed dynamically while ADSM is running. The capacity of the database can also be increased or decreased dynamically using server commands.

Mirroring provides the following advantages:

- Protects against media failures within the database or recovery log by providing up to three exact copies.
- Continues database operations without interruption if a database or recovery log volume fails by using a mirrored copy of the failed volume. The failed volume is placed offline and is automatically synchronized when brought back online.
- Provides a way to avoid costly database recoveries.
- Increases the database read performance by allowing the server to read data from any available volume copy.

Of course, mirroring requires additional DASD for the mirrored volumes.

The following scenario shows the importance of mirroring in the recovery process:

A sudden power outage occurs on the system on which an ADSM server is running. A partial page write occurs, thus corrupting the recovery log. The recovery log is now not completely readable. Without mirroring, transaction recovery operations cannot complete when the server is restarted. However, if the recovery log had been mirrored and a partial write is detected, a mirror volume can be used to construct valid images of the missing pages.

## Using Mirroring

Task	Required Privilege Class
Define database and recovery log volumes	System or unrestricted storage
Query mirrored volumes	Any administrator

This section explains how to:

- Allocate disk volumes to mirror the database and recovery log
- Define mirrored volume copies
- Monitor mirrored volume copies

### Allocating Volume Copies to Separate Physical Disks

By separating volume copies on different physical devices, you protect the server against media failure and increase the availability of the database and recovery log. You should place mirrored copies on separate devices, preferably behind separate controllers (if applicable).

**Note:** Mirrored volumes should have at least the same capacity as the original volumes.

### Defining Database or Recovery Log Mirrored Volumes

To ensure that the entire database or recovery log is mirrored, define a volume copy for each volume in the database or recovery log. For example, if the database consists of five volumes named VOL1, VOL2, VOL3, VOL4, and VOL5, then you must define five volume copies to mirror the database.

Figure 60 shows a mirrored database. In this example, VOL3 and VOLC are a group of mirrored volumes with the same portion of the database.




Figure 60. Mirrored Volumes

If you define volume copies from the command line and have multiple volumes to mirror, consider using a macro to issue more than one DEFINE command at a time. For example, to define a group of mirrored volumes for the database, use a macro to issue the following commands:

```
define dbcopy VOL1 VOLA
define dbcopy VOL2 VOLB
define dbcopy VOL3 VOLC
define dbcopy VOL4 VOLD
define dbcopy VOL5 VOLE
commit
```

After a volume copy is defined, ADSM synchronizes the volume copy with the original volume. This process can range from minutes to hours, depending on the size of the volumes and performance of your system. After synchronization is complete, the volume copies are mirror images of each other.

### Requesting Information about Mirrored Volumes

Any administrator can query the server to request information about database or recovery log volumes by using the QUERY DBVOLUME and QUERY LOGVOLUME commands. For example, to query the server for information about mirrored database volumes, enter:

```
query dbvolume
```

Figure 61 shows the database volume report you receive after issuing the QUERY DBVOLUME command. In this report, each vertical column displays an entire image of the database or recovery log. For example, VOLA, VOLB, VOLC, VOLD, and VOLE represent one image of the database. On the graphical user interface, each vertical column is called a *volume set*.

Each horizontal column displays a *group of mirrored volumes*. For example, VOL1, VOLA, and VOL600 represent three volume copies that are an exact image of one another.

Volume Name (Copy 1)	Copy Status	Volume Name (Copy 2)	Copy Status	Volume Name (Copy 3)	Copy Status
VOL1	Sync'd	VOLA	Sync'd	VOL600	Sync'd
VOL2	Sync'd	VOLB	Sync'd	VOL500	Stale
VOL3	Sync'd	VOLC	Sync'd	VOL400	Stale
VOL4	Sync'd	VOLD	Sync'd		Unde- fined
VOL5	Sync'd	VOLE	Sync'd		Unde- fined

Figure 61. Information about Database Volumes

When you query for database or recovery log information, the status of a mirrored volume can be:

<b>Sync'd</b>	The volume is the only copy or is synchronized with other volumes in the group.
<b>Stale</b>	The volume is not available for recovery because synchronization has not completed or begun.
<b>Offline</b>	A volume has been made unavailable to the server by varying the volume offline. See “Varying Disk Volumes Online or Offline” on page 55 for details.
<b>Undefined</b>	No volume has been defined as a mirrored copy.

### Controlling Database and Recovery Log Performance

You can control database and recovery log processing and performance by using two options in the server options file. The `MIRRORREAD` option and the `MIRRORWRITE` option specify modes for reading and writing database and recovery log pages. See *ADSM Installing the Server and Administrative Client* for details.

---

## Recovering by Using Mirrored Copies of the Database

If a volume fails because of media failure and mirroring has been in effect, you can recover the volume by taking the following steps:

1. View the current status on database and recovery log volumes and volume copies (`QUERY DBVOLUME` or `QUERY LOGVOLUME`).
2. If necessary, place the failing volume offline from ADSM (`DELETE DBVOLUME` or `DELETE LOGVOLUME`). The server usually does this automatically.
3. Fix the failing physical device.
4. Allocate space to be used for the new database or recovery log volume copy (`DSMFMT`).
5. Bring the ADSM volume online (`DEFINE DBCOPY` or `DEFINE LOGCOPY`).

After a database or recovery log volume copy is defined, the server synchronizes the volume copy with its associated database or recovery log volume.

---

## Database Backup and Recovery

You can take full or incremental backups of the database while the server is operational and available to clients. If a disaster occurs, you can use the backed up copies to restore the database. ADSM lets you recover the database to its most current state (roll-forward recovery) or to a specific point in time.

The best database protection is to mirror the database and recovery log and periodically back up the database. To take the fullest advantage of ADSM’s database backup and recovery feature, you must create volume history files (see “Establishing Volume History Backup Files” on page 270) and device configuration files (see “Establishing Device Configuration Backup Files” on page 271).

To ensure the fastest recovery time and highest availability of the database, both the database and recovery log should be mirrored. However, if you do not mirror the database and if the recovery log is available, you can still restore the database to its most current state. Therefore, whatever strategy you use, you should always

mirror the recovery log. Mirroring only the recovery log requires much less storage space than mirroring the database.

## Roll-Forward Recovery

With the recovery log available, roll-forward recovery lets you restore a database to its most current state. Together, mirroring and roll-forward recovery provide the most comprehensive protection for an ADSM database.

When enabling roll-forward recovery, consider the following:

- The recovery log must be large enough to store additional recovery log records.

Operating in roll-forward mode significantly increases recovery log storage requirements. The extent of the increase is determined by the number of database transactions since the last database backup. Database transactions are created by server operations including client backups and archives, storage pool migrations, tape reclamations, and expiration processing. In roll-forward mode, the recovery log tracks all transactions since the last database backup. In using point-in-time mode, only current, uncommitted transactions are kept in the recovery log.

- More frequent database backups help control recovery log size.

When a full or incremental database backup is completed, recovery log records preceding the backup are deleted, freeing up recovery log storage for reuse. Using the database backup trigger automatically starts database backups based on recovery log utilization.

You should monitor recovery log storage, recovery log usage, and the database backup trigger to achieve a balance that best meets your needs.

- Mirroring a larger recovery log also requires additional disk space.
- The frequency of required database backups depends on the size of the recovery log, how the database backup trigger is set, and the volume of ADSM transactions.
- Roll-forward mode is an alternative to database mirroring. It lets you recover to the most recent state following a media failure, without doubling your database storage requirements. However, unlike mirroring, roll-forward recovery does not provide continuous operations on media failure.
- Roll-forward mode is not a disaster recovery utility. When the recovery log is not available, as with a disaster or with a recovery log media failure when not using mirroring, you must recover to a point in time.

## Point-in-Time Recovery

Point-in-time recovery is normally used for exceptional situations such as disaster recovery or to remove the effects of errors that can cause inconsistencies in the database. You can also use point-in-time recovery if you do not want to create a larger recovery log size for ADSM. However, if the server supports HSM clients roll-forward recovery is preferable. Space-managed files should be protected as fully as possible from hardware failure. If you enable roll-forward recovery, point-in-time recovery is also available.

For disaster recovery purposes, you can run full or incremental backups as often as necessary to ensure that an ADSM database can be restored to an acceptable point in time. You can send the backups to an offsite location for safekeeping.

You can also choose to enable point-in-time recovery only. The disadvantages of enabling only point-in-time recovery are:

- You cannot restore the database to its most current state. You can only restore it to a time at which a backup was taken.
- Because restoring a single database volume requires roll-forward recovery, you must restore the entire data base even if only one volume is damaged.
- You should mirror the database. Mirrored copies of a database can consume a significant amount of disk space.

The advantage of using point-in-time, however, is that the recovery log can be defined with a much smaller size.

## Determining Which Recovery Mode to Run

There are several key tradeoffs in deciding whether to run in roll-forward mode or point-in-time mode.

Roll-forward mode provides the ability to recover the database to the most recent state upon a media failure on a database volume. Point-in-time (NORMAL) mode requires significantly less recovery log space.

For most situations, running in point-in-time mode is sufficient if you can afford to mirror the database and recovery log or place them on devices that guarantee availability. If you are mirroring the database, roll-forward mode only has value if all copies of a database volume fail concurrently and all recovery log volumes are still available.

If your database is so large that you cannot afford to mirror it, roll-forward mode provides a very good alternative. It allows you to recover to the time of failure. However, some of the storage you save from not mirroring the database will need to be reinvested in a much larger recovery log.

## Determining Which Type of Backup to Run

A full backup takes longer to run than an incremental backup because it copies the entire database. However, recovery time is faster with a full backup because only one set of volumes needs to be loaded to restore the entire database. A full backup is required under specific conditions, but you can run up to 32 incremental backups between each full backup.

An incremental backup takes less time to run because it copies only those database pages that have changed since the last time the database was backed up. However, incremental backups increase the time it takes to recover a database because a full backup must be loaded first, followed by some or all of the incremental backups in the same database backup series.

## When to Perform Backups

To ensure that you have the most recent database information, you should back up the database if:

- Significant client backup or archive activities have taken place on the server
- Migration has moved data between storage pools
- Reclamation has moved client files on sequential storage volumes

- Move data or delete volume commands have been used to rearrange data or remove volumes from a storage pool
- Storage pool backups have been done

Under typical conditions, you should probably back up your storage pools each day, immediately followed by a database backup. Depending on the amount of client data and frequency of the activities mentioned above, you may back up less often.

**Note:** In roll-forward mode with a database backup trigger defined, backups are automatic.

## Database Backup Trigger

When you are running in roll-forward mode, a database backup trigger causes ADSM to run a backup of the database automatically, based on the percentage of space used in the recovery log. When the amount of space occupied in the recovery log reaches the percentage that you specify, ADSM automatically runs a full or incremental backup of the database and deletes any unnecessary recovery log records.

**Note:** ADSM uses the database backup trigger only when roll-forward mode is in effect.

Setting a database backup trigger is optional, but it is recommended to ensure that the recovery log does not run out of space before the next backup is run.

## Volume History Backup Files

To perform a point-in-time restore, ADSM needs information about the volumes used for the backups. After a point-in-time restore, you need to know which sequential access storage pool volumes to audit.

The following information, which the ADSM server stores in the database, is needed for a point-in-time restore:

- Volumes used for database backups
- Volumes used to export administrator, node, policy, or server data
- Sequential access storage pool volumes that have been added, reused, or deleted

During a restore of the database, this information cannot be obtained from the database. Therefore, you should have at least one backup copy of the volume history information. You can specify one or more files in which the server backs up the volume history information. Whenever ADSM updates volume history information in the database, it also updates the volume history files. For details about setting up the volume history file, see “Establishing Volume History Backup Files” on page 270.

**Note:** You can recover the database without a volume history file. However, because you must examine every volume that may contain database backup information, this is a time consuming and error-prone task.



## Device Configuration Backup Files

When you define, update, or delete a device class, ADSM updates the database. To restore a database, ADSM requires a definition for the device class from which backup data is to be read. When the database is being restored no definitions can be read from the database. Therefore, you should have at least one backup copy of the device configuration information. You can specify one or more files in which the server backs up the device configuration information. Whenever ADSM updates device configuration information in the database, it also updates the information in the backup files. For details about setting up the device configuration file, see “Establishing Device Configuration Backup Files” on page 271.

During a database restore operation, ADSM tries to open the first device configuration file. If it cannot open or read that file successfully, ADSM tries to use any remaining device configuration files (in the order in which they appear in the DSMSEV.OPT file) until it finds one that is usable. If none can be found, you must recreate the file. See “Recreating a Device Configuration File” on page 272 for details.

---

## Using Database Backup Features

Task	Required Privilege Class
Set the recovery log mode	System or unrestricted storage
Extend or reduce the size of the recovery log	System or unrestricted storage
Define, update, or delete the database backup trigger	System or unrestricted storage
Query the database backup trigger	Any administrator
Back up or delete the volume history file or the device configuration file	System or unrestricted storage
Query the volume history file or the device configuration file	Any administrator
Backup the database	System or unrestricted storage

This section discusses how to set up and manage database backup. The following topics are included:

- Defining device classes for backup
- Establishing volume history backup files
- Establishing device configuration backup files
- Performing the database backups
- Setting the recovery log mode
- Setting the database backup trigger
- Adjusting the size of the recovery log

## Defining Device Classes for Backup

You can use existing device classes to use for backups or define new ones. For incremental backups you can specify a device class different from the one used for full backups.

For example, you can write full backups to a tape device, and incremental backups to a disk device. Specifying a device class with a device type of FILE is useful for incremental backups run automatically by ADSM based on the setting specified with the DEFINE DBBACKUPTRIGGER command.

You should specify the same device class for all incremental backups that apply to the same full backup to make restoring those backups as fast as possible.

In addition, to avoid a situation in which a backup needs to be run based on the database backup trigger but no mount point is available, you can define one or more device classes to be used only for automatic backups. This process reserves one or more mount points specifically for this purpose. However, if you share the device class with other operations and because BACKUP DB is a high priority operation, it automatically cancel lower priority operations such as reclamation if all the mount points are in use. This frees a mount point for the database backup.

## Establishing Volume History Backup Files

Because volume history information is vital for a point-in-time recovery on a lost or damaged database, you should ensure that you have at least one backup copy of this information. You should put the volume history file on a disk other than the one on which your database volumes reside. You can also store a remote copy, for example, on an NFS-mounted file system.

You can specify one or more backup copies by defining them with the VOLUMEHISTORY option in the server options file (DSMSERV.OPT). Maintaining multiple backup copies of the volume history information in different locations can help to ensure that the information is available if it is needed. For disaster recovery purposes, you can also print the volume history information and store it offsite.

For example, to define the **volhist.bk1** and **volhist.bk2** volume history backup files, enter the following in DSMSERV.OPT:

```
volumehistory volhist.bk1
volumehistory volhist.bk2
```

**Note:** With the VOLUMEHISTORY option, you can specify files in different libraries.

After you have made these entries, halt and restart the server for them to become effective. When you halt the server, ensure that you are not affecting any client sessions. For information on how to stop the server, see “Halting the Server” on page 50.

After you have restarted the server, whenever ADSM updates information about volumes in the database, it also updates the same information in the files you have defined in the VOLUMEHISTORY option.

You can also use the BACKUP VOLHISTORY command at any time to store a backup copy of volume history information in one or more specified files.

For example, to back up the volume history information in the **volhist.bk3** and **volhist.bk4** files, enter:

```
backup volhistory filenames=volhist.bk3,volhist.bk4
```

If you do not specify files in the command, ADSM backs up the volume history information in *all* files specified with the VOLUMEHISTORY option in the DSMSEV.OPT file.

You can use the QUERY VOLHISTORY command to display volume history information. To display all volume history information up to yesterday, enter:

```
query volhistory enddate=today-1
```

You can delete volume history information by using the DELETE VOLHISTORY command.

For example, to delete information that is 30 days older or older, enter:

```
delete volhistory todote=today-30
```

**Note:** When volumes are deleted from the volume history with the DELETE VOLHISTORY command, they return to scratch status in the libraries attached to the server and may be reused.

You should specify volume history files in the server options file or schedule the BACKUP VOLHISTORY command to follow your scheduled database backups. A copy of your most recent volume history file should be stored offsite with your database backups.

## Establishing Device Configuration Backup Files

Because device configuration information is vital to restoring a lost or damaged database, you should ensure that you have at least one backup copy of this information. You should specify one or more backup copies by defining them with the DEVCONFIG option in the server options file (DSMSEV.OPT). When you define, update, or delete a device class, ADSM writes backup copies of the resulting information in the files specified with the DEVCONFIG option.

Maintaining multiple backup copies of device configuration information in different locations can help ensure that the information is available if it is needed. For disaster recovery purposes, you can also print the device configuration information and store it offsite.

For example, to define the **devconf.bk1** and **devconf.bk2** device configuration backup files, enter the following in the DSMSEV.OPT file:

```
devconfig devconf.bk1  
devconfig devconf.bk2
```

After you have made these entries, halt and restart the server for them to become effective. When you halt the server, ensure that you are not impacting any client sessions. For information on how to stop the server, refer to “Halting the Server” on page 50.

After you have restarted the server, whenever you use the DEFINE, UPDATE, or DELETE DEVCLASS commands, ADSM backs up copies of the resulting definitions to the files that you have defined in the DEVCONFIG option.

You can also use the BACKUP DEVCONFIG command at any time to store a backup copy of device configuration information in one or more specified files.

For example, to back up the device configuration information in the **devconf.bk3** and **devconf.bk4** files, enter:

```
backup devconfig filenames=devconf.bk3,devconf.bk4
```

If you do not specify files in the command, ADSM stores copies of the device configuration file in *all* files specified with the DEVCONFIG option in the DSMSERV.OPT file. A copy of your most recent device configuration file should be stored offsite with your volume history file and database backups.

If you lose your device configuration file and need it to restore the database, you must recreate it manually. See “Recreating a Device Configuration File” for details.

## Recreating a Device Configuration File

The following commands read and execute the device configuration file:

- DSMSERV RESTORE DB
- DSMSERV LOADDB
- DSMSERV DISPLAY DBBACKUPVOLUME

If no device configuration file is found, you must recreate it before you can start the restore operation. The file consists of DEFINE DEVICE commands that ADSM runs when you issue one of the previous DSMSERV commands. You need to provide only those definitions needed to mount the volumes read by the DSMSERV command. If you are restoring or loading from a FILE device class, you will need only the DEFINE DEVCLASS command. The device configuration file must follow these conventions:

- You can use command defaults.
- The file can include blank lines.
- A single line can be up to 1550 characters.
- The file can include continuation characters and comments as described in the *ADSM Administrator's Reference*.

The following figure shows an example of a device configuration file:

```
/* IBM AdStar Distributed Storage Manager Device Configuration */  
define devclass cart devtype=cartridge
```

## Performing the Database Backups

Before using the database backup and recovery feature, you must first do a full back up of your database.

For example, to perform a full backup of your database to the TAPE1 device class, enter:

```
backup db type=full devclass=tape1
```

In this example, ADSM writes the backup data to scratch volumes. You can also specify volumes by name. After you have done a full backup, you can perform incremental backups, which copy only the changes to the database since the previous backup.

To do an incremental backup of the database to the TAPE1 device class, enter:

```
backup db type=incremental devclass=tape1
```

## Setting the Database Backup Trigger

The DEFINE DBBACKUPTRIGGER command specifies a percentage of used space in the recovery log. When that percentage is reached and if the log mode is set to ROLLFORWARD, ADSM automatically runs a full or incremental database backup database and deletes any unnecessary recovery log records.

Setting a database backup trigger is optional, but recommended to ensure that the recovery log does not run out of space before the next backup is run.

**Note:** If the log mode is set from NORMAL to ROLLFORWARD, the next database backup must be a FULL backup. The server does not start saving log records for roll-forward recovery until this full backup completes successfully. If a database backup trigger is defined when you set the log mode to ROLLFORWARD, the full backup is done automatically.

To set the database backup trigger at 60 percent for an automatic backup of the database, enter:

```
define dbbackuptrigger logfullpct=60 devclass=TAPE1
```

If you do not specify the LOGFULLPCT parameter in the command, the backup trigger defaults to 50 percent. In addition, ADSM runs 6 incremental backups to every full backup unless you specify a different value in the command.

For example, to run 20 incremental backups to every full backup, enter:

```
define dbbackuptrigger logfullpct=60 devclass=tape1 numincremental=20
```

Each incremental backup, whether run automatically by ADSM or run by issuing the BACKUP DB command, is added to the count of incremental backups run.

Each full backup, whether run automatically by ADSM or run by issuing the BACKUP DB command, resets the count for incremental backups to zero. When you specify 0 for the NUMINCREMENTAL parameter, ADSM automatically runs only full backups.

After you initially set the database backup trigger, you might find that automatic backups are being run too often or not enough. Before taking any action, check the current settings in order to understand how much to change the backup trigger percentage.

To check your current settings, enter:

```
query dbbackuptrigger
```

ADSM displays the following information:

```
Backup Trigger Status: Enabled
Device Class for Full Backups: FULL
Device Class for Incremental Backups: INCR
Log Full Percentage: 60
Number of Incremental Backups per Backup Series: 7
Incremental Backups taken in this Backup Series: 0
Next Triggered Backup Type: Incremental
Database Backup in Progress?: No
Database Backup Type:
Last Update by (administrator): HINEY
Last Update Date/Time: 03/06/1995 10:49:23
```

For example, assume that the database backup trigger is automatically running backups more often than you want. The above information shows that the Log Full Percentage attribute is set at 60 percent. To increase that attribute to 70 percent, enter:

```
update dbbackuptrigger logfullpct=70
```

If you no longer want to use the database backup trigger, enter:

```
delete dbbackuptrigger
```

When you delete the database backup trigger, ADSM does not run backups of the ADSM database automatically and does not remove unnecessary log records from the recovery log.

After you delete the database backup trigger, be sure to use the SET LOGMODE command to change the log mode to NORMAL. If the database backup trigger

automatically runs backups more often than you want and the setting is high (for example, 90%), you should probably increase the recovery log size.

**Note:** If you turn off the trigger and stay in roll-forward mode, transactions fail when the log fills. If you set the log mode to NORMAL, you cannot perform roll-forward recovery.

## Setting the Recovery Log Mode

The recovery log mode determines how long ADSM saves records in the recovery log. You use the SET LOGMODE command to specify whether or not to save enough log records to perform roll-forward recovery.

The default log mode is NORMAL, which does not allow for roll-forward recovery. ADSM saves only those records required to restore the database to the point it was at when the last backup was run.

The advantage to using NORMAL log mode is that it requires less storage space for the recovery log than running with roll-forward log mode requires.

To set the log mode to normal, enter:

```
set logmode normal
```

To make roll-forward recovery available, you must set the logmode to ROLLFORWARD. When you set the logmode to ROLLFORWARD, ADSM saves all recovery log records that reflect changes to the database since the last time a database backup was run. ADSM saves the recovery log records needed to roll the database forward to its most current state after loading the most recent backup series. Each time the database is backed up, ADSM deletes any recovery log records it no longer needs for that purpose.

To set the log mode to ROLLFORWARD, issue the following command:

```
set logmode rollforward
```

## Adjusting the Size of the Recovery Log

The frequency of automatic backups depends not only on the percentage set for the database backup trigger, but also on the size of the recovery log and the volume of ADSM transactions. The considerations for adjusting the size of the recovery log differ depending on the log mode that you select.

If the log mode is set to NORMAL, adjust the size of the recovery log based only on the volume of concurrent ADSM transactions. As more clients are added and the volume of concurrent transactions increases, you can extend the size of the log.

In ROLLFORWARD mode, the size of the recovery log, the percentage set for the database backup trigger, and the volume of ADSM transactions all affect the frequency at which backups are required.

Generally, extending the size of the recovery log decreases the frequency of automatic backups, and reducing the size of the recovery log increases the frequency. However, after the size of the recovery log is established, you can also increase or decrease the frequency of automatic backups by adjusting the percentage specified for the database backup trigger. In addition, fluctuations in the volume of ADSM transactions can also affect the frequency of automatic backups.

To increase the size of the recovery log, use the EXTEND LOG command. For example, to increase the size of the recovery log by 100MB, enter:

```
extend log 100
```

**Note:** You can extend the log only if there is space. The QUERY LOG command displays the maximum extension permitted. If there is no space for extension, you must add volumes to the recovery log before extending it.

---

## Recovering by Using Backed Up Copies of the Database

This section explains how to recover from a disaster by using backed up copies of the database and recovery log. The following topics are included:

- Restoring a database by using point-in-time recovery
- Restoring a database to its most current state

To perform a restore, you should have the following information, preferably stored offsite:

- Back up volumes of the database
- Copy storage pool volumes
- On the same type of tape volumes as the backups or on diskette, or as printouts:
  - The server options file (DSMSERV.OPT)
  - Volume history file
  - Device configuration file
  - Output from QUERY DBVOLUME FORMAT=DETAILED and QUERY LOGVOLUME FORMAT=DETAILED (for details of database and recovery log setup)

## Restoring a Database by Using Point-in-Time Recovery

You need backup copies of volume history and device configuration information for the recovery operation. See “Volume History Backup Files” on page 268 and “Device Configuration Backup Files” on page 269 for details about defining these files.

If the volume history file is not available, you must mount tape volumes in the correct order or specify their order on the DSMSERV RESTORE DB command.

**Note:** After you restore the database, you lose any volume history information pointed to by the options file. You will need this information to identify the



volumes to be audited. Therefore, before issuing the DSMSEV RESTORE DB command, rename and save a copy of the volume history file.

If you lose the device configuration file, you can restore it manually (see “Recreating a Device Configuration File” on page 272).

To restore the database to a specific point in time, use the DSMSEV RESTORE DB command. ADSM restores the database as follows:

- Reads the volume history file to locate the last full backup that occurred on or before the specified date and time.
- Using the device configuration file, requests a mount of the first volume, which should contain the beginning of the full backup.
- Restores the backup data from the first volume.
- Continues to request mounts and to restore data from the backup volumes that contain the full backup and any incremental backups that occurred on or before the date specified.

For example, to restore the database to a backup series that was created on January 3, 1995, enter:

```
dsmserv restore db todate=01/03/95
```

From the old volume history information, you need a list of all the volumes that were reused, added, and deleted since the original backup. Using this list, perform the following steps:

1. Audit all disk volumes, all reused volumes (STGREUSE), and any deleted volumes that you could locate. This process identifies files recorded in the database that can no longer be found on the volume. If a copy of the file is in a copy storage pool, the file on the audited volume is marked as damaged. Otherwise, the file is deleted from the database and is lost. If the audit detects any damaged files, issue the RESTORE STGPOOL command to restore those files after you have audited all the volumes in that storage pool.
2. Mark as destroyed any volumes that cannot be located, and recover those volumes from copy storage pool backups. If no backups are available, delete the volumes from the database by using DISCARDATA=YES.
3. Re-add any volumes that were added since the backup.

Some files may be lost if they were moved since the backup (due to migration, reclamation, or move data requests) and ADSM cannot locate them. You can minimize this loss by using the REUSEDELAY parameter when defining or updating storage pools. This parameter delays volumes from being returned to scratch or being reused.

If you have backed up your storage pools using the BACKUP STGPOOL command, instead of auditing volumes, you can update the reused volumes (STGREUSE) and the deleted volumes as DESTROYED. Then restore the storage pools. By backing up your storage pool data in addition to backing up your database, you reduce the risk of losing data. To further minimize loss of data, you can:

- Back up the database immediately after you back up the storage pools.

- Turn off migration and reclamation while you back up the database.
- Do not perform any MOVE DATA operations during your back up of the database.
- Use the REUSEDELAY interval to prevent your copy storage pool volumes from being reused or deleted before they might be needed.
- Use offsite copy storage pool volumes. If, after backing up the storage pools, you mark as OFFSITE all of the volumes in your copy storage pool, they are preserved and are not reused or mounted until they are brought onsite. Ensure that you mark the volumes as offsite before you back up the database.

If your old volume history file shows that any of the copy storage pool volumes needed to restore your storage pools have been reused (STGREUSE) or deleted, you may not be able to restore all your files, and you will receive message during your restore operation.

After a restore, the volume inventories for ADSM and for your tape management system may be inconsistent. You will have to manually synchronize these inventories.

## Restoring a Database to its Most Current State

You can use roll-forward recovery to restore a database to its most current state if:

- The log mode was set to ROLLFORWARD continuously from the time the last full backup was created until the time the database was damaged or lost.
- At least one mirror copy of the recovery log is available and all recovery log volumes are intact.

For roll-forward recovery, ADSM uses the last backup series created for the database. A backup series includes a full backup and any incremental backups that apply to that full backup, plus all recovery log records that reflect changes to the database since the last backup in the series was run. To restore the database to its most current state, use the DSMSEV RESTORE DB command.

For example, to restore the database to its most current state, enter:

```
dsmserv restore db
```

The following example shows the importance of storage pool backups with a point-in-time restore. In this example, caching is not active, and the storage pool was not backed up with the BACKUP STGPOOL command.

*9:30am* Client A backs up its data to Volume 1.

*Noon* The system administrator backs up the database.

*1:30pm* Volume 1, containing Client A's files, is migrated to tape (Volume 2).

*3:00pm* Client B backs up its data to Volume 1.

The server places Client B's files in the location that contained Client A's files prior to the migration.

*3:30pm* The server goes down.

3:40pm The system administrator reloads the noon version of the database by using the DSMSERV RESTORE DB command.

4:40pm Volume 1 is audited. The following then occurs:

1. The server compares the information on Volume 1 and Volume 2 with the restored database (which matches the database at noon).
2. The audit does not find Client A's files on Volume 1 where the reloaded database indicates they should be. Therefore, the server deletes these Client A file references.
3. The database has no record that Client A's files are on Volume 2, and the files are, in effect, lost.
4. The database has no record that Client B's files are on Volume 1, and the files are, in effect, lost.

If roll-forward recovery had been used in the previous scenario, the database would have been rolled forward to 3:30pm when the server went down, and neither Client A's files nor Client B's files would have been lost. If a point-in-time restore of the data base had been performed and the storage pool had been backed up, Client A's files would not have been deleted by the volume audit and could have been restored with a RESTORE VOLUME or RESTORE STGPOOL command. See Chapter 14, "Recovering Data" on page 261 for information about recovery modes.

---

## Storage Pool Backup and Recovery

You can create backup copies of client files that are stored in your primary backup, archive, and space management pools. The backup copies are stored in *copy storage pools* that can be used to restore the original files in case they become damaged, lost, or unusable.

A typical storage hierarchy migrates from disk to tape. These primary storage pools should be backed up incrementally to the same copy storage pool each day. By backing up to the same copy storage pool, you ensure that files are not recopied as they migrate to the next storage pool.

With scheduled storage pool backups and migrations and with sufficient disk storage, most copies can be made from the disk storage pool before the files are migrated to tape, thus avoiding unnecessary mounts.

Backing up storage pools introduces additional space requirements on the database. ADSM maintains control information concerning the location, name, and characteristics of copy storage pool files in the server database. A copy storage pool file does not require an additional database entry. Instead, it adds to the existing database entry for the file information about the copy pool location. The information kept for each file copy requires about 200 bytes of space. In addition, a small amount of space is needed for internal database indexing. As more files are added to the copy storage pools, reevaluate your database size requirements.

---

## Using Storage Pool Backup Features

Task	Required Privilege Class
Define, back up, or restore storage pools	System, unrestricted storage, or restricted storage (only for those pools to which you authorized)
Update volumes	System or operator
Restore volumes	System, unrestricted storage, or restricted storage (only for those pools to which you authorized)
Query volumes or storage pools	Any administrator

ADSM provides storage pool backup features that can prepare you for an orderly and successful recovery if a disaster or media loss occurs. With ADSM you can perform incremental backups of primary storage pools. Incremental backups of storage pools, along with database backups and periodic reclamation of offsite volumes, should provide protection to meet most recovery requirements. However, you may sometimes want to do full backups of primary storage pools for the following reasons:

- A concern for the shelf life of the backup media
- Legal or audit requirements that specify how backups should be done
- A preference for the full backup and applied incremental approach

You can do full backups of storage pools by backing up to a new copy storage pool. Additional backups to this new copy pool would be incremental. A full backup process can be set up using this approach.

For example, you want to take full backups weekly and incremental backups daily, and you want to keep four weeks' worth of backups. You can do this by backing up to a new copy storage pool each week. When a copy pool is 4 weeks old, you can delete all the volumes in that copy pool (using the `DELETE VOLUME` command with `DISCARDATA=YES`). You could then reuse the copy pool and its volumes for the next weeks' backups. You would also set the `REUSEDELAY` value for the copy storage pool to 28 days.

There are some drawbacks to this approach:

- This approach maintains five copies of each file, the primary copy and the four backups. Information on each copy must be kept in the database. The backup copies do not require full database entries (as do the primary copies), but they do take up database space.

---

## Recovering by Using Backed Up Copies of Storage Pools

By using backed up copies of storage pools, you can recover from a disaster or media loss. "Recovery from a Disaster" on page 282 and "Recovery from Media Loss" on page 284 present scenarios for these situations.

---

## ADSM Backup and Recovery Scenarios

This section presents scenarios for protecting and recovering an ADSM server. You may want to modify the procedures to meet your requirements.

In these scenarios, a company's storage hierarchy consists of the default BACKUPPOOL, ARCHIVEPOOL, and SPACEMGPOOL storage pools. All the storage pools use random access media. The server migrates the files in these storage pools to the TAPEPOOL storage pool, which uses cartridge media.

### Backup

The company's standard procedures include the following:

- Perform a full backup of the database once a week and incremental backups on the other days.
- Perform reclamation of its copy storage pool, once a week. Reclamation for the copy storage pools is turned off at other times.
- Back up its storage pools every night and the server database immediately after the storage pool backups.
- Ship the database and data storage backups to an offsite location every day.

To prepare for a possible disaster, the administrator does the following:

1. Creates a copy storage pool named DISASTER-RECOVERY. Only scratch cartridges are used, and the maximum number of scratch volumes is set to 100. The copy storage pool is defined by entering:

```
define stgpool disaster-recovery cartridge pooltype=copy maxscratch=100
```

2. Defines schedules to do the following operations every day:

- a. Incremental backups of the primary storage pools are done each night by issuing the following commands:

```
backup stgpool backuppool disaster-recovery maxprocess=4  
backup stgpool archivepool disaster-recovery maxprocess=4  
backup stgpool spacemgpool disaster-recovery maxprocess=4  
backup stgpool tapepool disaster-recovery maxprocess=4
```

These commands use multiple, parallel processes to perform an incremental backup of each primary storage pool to the copy pool. The only files that are backed up to the DISASTER-RECOVERY pool are those files for which a copy does not already exist in the copy storage pool.

- b. Select volumes in the DISASTER-RECOVERY copy storage pool that have read-write or read-only access and which are onsite, at least partially filled, and change the volume's access mode to OFFSITE by entering:

```
update volume * access=offsite location='vault site info' -  
wherestgpool=disaster-recovery whereaccess=readwrite,readonly -  
wherestatus=filling,full
```

- c. Backup the database by using the BACKUP DB command:

```
backup db type=incremental devclass=devclassname scratch=yes
```

3. Performs the following operations each night after the scheduled operations previously discussed have completed:
  - a. Back up the volume history, device configuration, and server options files to facilitate recovery.
  - b. Move the volumes marked offsite and the database backup volumes, volume history files, device configuration files, and server options files (all produced in the previous steps) to the offsite location.
  - c. Identify the offsite volumes that are ready to be returned to the onsite location by using the QUERY VOLUME command:

```
query volume stgpool=disaster-recovery access=offsite status=empty
```

These volumes, which have become empty, through expiration, reclamation, and file space deletion, have waited the delay time specified by the REUSEDELAY parameter. The administrator periodically returns outdated backup database volumes. These volumes are displayed with the QUERY VOLHISTORY command and can be released for reuse with the DELETE VOLHISTORY command.

4. Return the volumes identified in step 3c to the onsite location and update their access to read-write. The REUSEDELAY setting that the company uses for the DISASTER-RECOVERY copy storage pool is set to a large enough value to ensure that no volume is returned to scratch that might be needed if a “non-outdated” database is restored. The volumes that have been empty for the number of days specified for REUSEDELAY are then automatically returned to scratch.

## Recovery from a Disaster

If a company makes the preparations described in “Backup” on page 281 it can recover from a disaster by using ADSM features.

In this scenario, the company’s database and all of its onsite data storage volumes are destroyed by fire. The processor with ADSM installed is not damaged. An administrator restores the server to the state that existed when the last backup was performed by doing the following:

1. Move the latest backup and all of the DISASTER-RECOVERY volumes from the offsite location to the onsite recovery location.
2. If a current, undamaged volume history file exists, save it.

3. Restore the volume history, device configuration, and server options files.
4. Restore the database from the latest backup level by issuing the DSMSEV RESTORE DB command, as described in “Recovering by Using Backed Up Copies of the Database” on page 276.
5. Change the access mode of all the existing primary storage pool volumes in the damaged storage pools to destroyed by issuing the UPDATE VOLUME command:

```
update volume * access=destroyed wherestgpool=backuppool
update volume * access=destroyed wherestgpool=archivepool
update volume * access=destroyed wherestgpool=spacemgpool
update volume * access=destroyed wherestgpool=tapepool
```

6. Identify if any volumes in the DISASTER-RECOVERY storage pool that are currently onsite by using the QUERY VOLUME command. If any are reported, these volumes were destroyed in the disaster and cannot be used for restore processing. Delete each of these copy storage pool volumes from the database by using the DELETE VOLUME command with the DISCARDATA option. Any files backed up to these volumes will not be restored.
7. Change the access mode of the remaining copy storage pool volumes in the DISASTER-RECOVERY pool to READWRITE by using the UPDATE VOLUME command:

```
update volume * access=readwrite wherestgpool=disaster-recovery
```

**Note:** Users can get files from ADSM at this point. If a user tries to get a file that was stored on a destroyed volume, the retrieval request goes to the copy storage pool. In this way, clients can be recovered with no data movement at all. When you update volumes brought from offsite to change their access, you greatly speed recovery time.

8. Define new volumes in the primary storage pool so the files on the damaged volumes can be restored to the new volumes. The new volumes also allow the end users to backup, archive, or migrate files to the server. You do not need to perform this step if you use only scratch volumes in the storage pool.
9. Recreate files in the primary storage pool from the copies located in the DISASTER-RECOVERY pool by entering:

```
restore stgpool backuppool maxprocess=4
restore stgpool archivepool maxprocess=4
restore stgpool spacemgpool maxprocess=4
restore stgpool tapepool maxprocess=4
```

These commands use multiple parallel processes to restore files to primary storage pools. After all the files have been restored for a destroyed volume, that volume is automatically deleted from the database.

10. To ensure against another loss of data, immediately back up all storage volumes and the database. Then resume normal activity, including weekly disaster backups and movement of data to the offsite location.

## Recovery from Media Loss

If a company makes the preparations described in “Backup” on page 281 it can recover from a disaster by using ADSM features.

In this scenario, an operator inadvertently destroys a tape volume (dsm087) belonging to the TAPEPOOL storage pool. An administrator performs the following actions to recover the data stored on the destroyed volume by using the offsite copy storage pool:

1. Change the access mode of tape volume dsm087 to UNAVAILABLE by using the UPDATE VOLUME command.
2. Determine the copy pool volumes that contain the backup copies of the files that were stored on the volume that was destroyed by entering:

```
restore volume dsm087 preview=volumesonly
```

This command produces a list of offsite volumes that contain the backed up copies of the files that were on tape volume dsm087.

3. Set the access mode of the copy volumes identified to UNAVAILABLE.

**Note:** This precaution prevents the movement of files stored on these volumes until volume dsm087 is restored.

4. Bring the identified volumes to the onsite location and set their access mode to READWRITE.
5. Restore the destroyed files by entering:

```
restore volume dsm087
```

This command sets the access mode of the dsm087 to DESTROYED and restores all of the files that were stored dsm087. The files are not actually restored to volume dsm087, but to another volume in the TAPEPOOL storage pool. All references to the files on dsm087 are deleted from the database and the volume itself is deleted from the database.

6. Set the access mode of the volumes used to restore dsm087 to OFFSITE using the UPDATE VOLUME command.
7. Return the volumes to the offsite location.

---

## Database Salvage Utilities

In the unlikely event of a seriously corrupted database and when no database backups are available, you can use the ADSM stand-alone database salvage utilities to try to recover. These utilities include the following three components:

- **DSMSERV DUMPDB**, which retrieves as many logical database records as possible into a database dump.



**Note:** The DSMSERV DUMPDB process does not access the recovery log. Any uncommitted transactions held in the log database are lost.

- **DSMSERV LOADDB**, which loads the dumped database to newly installed server. A badly corrupted database must be salvaged by reinitializing the server before LOADDB can be performed. Use the **DSMSERV INSTALL** program to initialize a new set of database and recovery log volumes.

You should allocate a new database and recovery log rather than destroy the damaged set, just in case the dump and load process needs to be repeated.

Storage pool volumes should not be reinitialized. Information in the salvaged database refers to client files stored on these volumes.

Ensure that the total size of the new database and log volumes are at least as large as the old ones.

- **DSMSERV AUDITDB**, which ensures that the database is returned to a consistent state. You should run AUDITDB with FIX=YES to recover the database. You must also set the recovery log mode to NORMAL or have just completed a database backup. Otherwise, the server may run out of recovery log space during the operation.

After the database audit completes and any errors are corrected, the ADSM server can be restarted. Its configuration should be identical to that just prior to the failure. The only exception could be transactions being processed at the time of failure. Any uncommitted are lost. You should review client activity at the time of failure. Updates to database pages maintained in the recovery log from the database buffer pool are also lost.

## Examples of Recovering a Corrupted Database Using DSMSERV Salvage Utilities

If the server database volumes are still readable, you can try to recover the existing database before reloading from a dump taken previously with the DSMSERV DUMPDB utility. In most cases, the information in existing database volumes is the most current.

### Recovery without a Previous Dump

To recover from existing database volumes, do the following:

1. Halt the server.
2. Ensure that a device configuration file is available and that the DEVCONFIG option is specified for this file in the server options file. If no device configuration file is available, create one manually (see "Recreating a Device Configuration File" on page 272 for details).
3. Dump the database by using the DSMSERV DUMPDB command.
4. Reinitialize database and recovery log by using the DSMSERV INSTALL command.
5. Reload the database by using the DSMSERV LOADDB command.
6. Audit the database by using the DSMSERV AUDITDB command and correct any problems.
7. Start the server.
8. Synchronize the database with volumes in the storage pool by using the AUDIT VOLUME.

## Recovery Using a Previous Dump

To recover from a database copy created by using the DSMSEV DUMPDB command, do the following:

1. Reinitialize the database and recovery log volumes by using the DSMSEV INSTALL command.
2. Ensure that a device configuration file is available and that the DEVCONFIG option is specified for this file in the server options file. If no device configuration file is available, create one manually (see "Recreating a Device Configuration File" on page 272 for details).
3. Reload the database by using the DSMSEV LOADDB command.
4. Audit the database by using the DSMSEV AUDITDB command and correct any problems.
5. Start the server.
6. Synchronize the database with volumes in the storage pool by using the AUDIT VOLUME.

---

# Glossary

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

This glossary may include terms and definitions from:

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

## A

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

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

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

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

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

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

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

**active version.** The most recent backup copy of a file stored by ADSM. Such a file is exempt from deletion until a backup detects that the user has either replaced

the file with a newer version, or has explicitly deleted the file from the workstation. Contrast with *inactive version*.

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

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

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

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

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

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

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

**ADSM.** ADSTAR Distributed Storage Manager.

**ADSM application programming 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 programming 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 data storage.

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

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

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

## B

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

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

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

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

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

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

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

**backup series.** See *database backup series*.

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

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

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

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

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

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

## C

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

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

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

**cartridge system tape (CST).** The base tape cartridge media used with 3480 or 3490 Magnetic Tape Subsys-

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

**CPI-C.** Common Programming Interface Communications.

**CST.** Cartridge system tape.

**CUA.** Common User Access.

## D

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

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

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

**DASD.** Direct access storage device.

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

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

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

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

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

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

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

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

**dataserver.** See *Tape Library Dataserver*.

**data set.** See *linear data set*.

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

**DDM.** Distributed Data Management.

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

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

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

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

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

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

**desktop.** On-screen representation of a desk top.

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

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

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

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

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

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

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

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

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

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

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

**DLL.** Dynamic link library.

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

**DOS.** Disk operating system.

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

**dsm.opt file.** See *client 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, and archive copies. 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's data storage.



ADSM does not necessarily store all the files from a single file space together, but can identify all the files in its data storage that came from a single file space.

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

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

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

**FTP.** File Transfer Protocol.

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

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

## G

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

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

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

**GUI.** Graphical user interface.

## H

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

**HDA.** Head-disk assembly.

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

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

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

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

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

**HSM client.** Hierarchical storage management client.

## I

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

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

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

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

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

include-exclude file (for UNIX clients) or in the client options file (for other clients).

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

**incremental backup.** (1) A function that allows users to 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 *mini-disk*.

**load.** See *mount*.

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

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

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

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

## M

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

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

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

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

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

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

**MB.** Megabyte.

**megabyte (MB).** (1) For processor storage and real and virtual memory, 2<sup>20</sup> or 1 048 576 bytes. (2) For disk storage capacity and transmission rates, 1 000 000 bytes.

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

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

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

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

**mm.** Millimeter.

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

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

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

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

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

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

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

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

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

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

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

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

**MVS.** Multiple Virtual Storage.

## N

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

**NETBIOS.** Network Basic Input/Output System.

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

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

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

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

**NFS.** Network File System.

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

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

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

## O

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

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

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

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

**Operating System/2 (OS/2).** An operating system used in IBM PC AT, PS/2, and compatible computers. OS/2 is one of the supported client environments and one of the supported server environments.

**operator privilege class.** An administrative privilege class that allows an administrator to issue commands that control the operation of the server. This privilege

class allows disabling or halting the server to perform maintenance, enabling the server, canceling server processes, and managing tape.

**optical disk.** A disk that contains data readable by optical techniques.

**optical drive.** A drive mechanism that rotates an optical disc.

**optical library.** A disk storage device that houses optical disk drives and optical disks, and contains a mechanism for moving optical disks between a storage area and optical disk drives.

**OS/2.** Operating System/2.

**OS/400.** Operating System/400.

**owner.** The owner of backup-archive files sent from a multiuser client node, such as AIX.

## P

**page.** (1) A block of instructions, data, or both. (2) In ADSM, a unit of space allocation within database volumes. (3) In a virtual storage system, a fixed block that has a virtual address and is transferred as a unit between real storage and auxiliary storage.

**paging.** (1) The action of transferring instructions, data, or both, between real storage and external page storage. (2) Moving data between memory and a mass storage device as the data is needed.

**pattern-matching expression.** A string expression that uses wildcard characters to specify one or more ADSM objects. See also *wildcard character*.

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

**platform.** The operating system environment in which a program runs.

**policy definition.** Server control information that includes information about policy domains, policy sets (including the ACTIVE policy set), management classes (including the default management class), copy groups, schedules, and associations between client nodes and schedules. A policy definition can be exported from a source server so that it can be imported to a target server at a later date.

**policy domain.** A policy object that contains policy sets, management classes, and copy groups that is used by a group of client nodes. See *policy set*, *management class*, and *copy group*.

**policy privilege class.** An administrative privilege class that allows an administrator to manage policy objects, register client nodes, and schedule client operations (such as backup services) for client nodes. Administrators can be authorized with unrestricted or restricted policy privilege. See *unrestricted policy privilege* or *restricted policy privilege*.

**policy set.** A policy object that contains a group of management class definitions that exist for a policy domain. At any one time there can be many policy sets within a policy domain but only one policy set can be active. See *management class* and *active policy set*.

**primary storage pool.** A named set of volumes that ADSM uses to store backup versions of files, archive copies of files, and files migrated from client nodes via ADSM space management. A primary storage pool may be backed up to a copy storage pool either automatically or by command. See *destination* and *copy storage pool*.

**privilege class.** A level of authority granted to an ADSM administrator. ADSM has five privilege classes: system, policy, storage, operator, and analyst. The privilege class determines which ADSM administrative tasks the administrator can perform. For example, an administrator with system privilege class can perform any administrative task.

**programmable workstation communication services (PWSCS).** A product that provides transparent high performance communications between programs running on workstations or on host systems.

**protection status.** A device class attribute that specifies whether to update the RACF profile to identify which users have access to cartridge tapes associated with this device class on MVS servers.

**PWSCS.** Programmable workstation communication services.

## Q

**QIC.** Quarter-inch cartridge (a type of magnetic tape media).

## R

**random access media.** Any volume accessed in a nonsequential manner. In ADSM, volumes are accessed in a nonsequential manner if they reside in the DISK device class.

**randomization.** The percentage of the startup window that the server can use to randomize start times for different client nodes associated with a schedule.

**rebinding.** The process of associating a file with a new management class name. For example, rebinding occurs when the management class associated with a file is deleted. See *binding*.

**recall.** A function that allows users to access files that have been migrated from their workstations to ADSM storage via ADSM space management. Contrast with *migrate*.

**reclamation.** A process of consolidating the remaining data from many sequential access media onto a single new sequential access media.

**reclamation threshold.** A value that specifies a percentage of space on sequential access media volumes that can be occupied by reclaimable space. The remainder of the space is for active data. (Space becomes reclaimable when files are expired.)

**recovery log.** A log of updates that are about to be written to the database. The log can be used to recover from system and media failures.

**recovery log buffer pool.** Used to hold new transactions records until they can be written to the recovery log.

**reduce.** The process of freeing up enough space to allow you to delete a volume from the database or recovery log. Contrast with *extend*.

**REEL.** On a VM server, a device class that is defined by ADSM at installation. It is used with VM servers to categorize tape devices that support tape reels, such as the 3420 9-track tape device.

**register.** Defines a client node or administrator who can access the server. See *registration*.

**registration.** The process of identifying a client node or administrator to the server.

**reply operator.** On a VM server, a VM user ID that specifies an operator who will reply to tape mount requests by the server.

**restore.** The process of returning a backup copy to an active storage location for use. ADSM has processes for restoring its database, storage pools, storage pool volumes, and users' backed-up files. For example, users can copy a backup version of a file from the storage pool to the workstation. The backup version in the storage pool is not affected. Contrast with *backup*.

**restricted policy privilege.** An administrative privilege class that enables an administrator to manage policy objects only for the policy domains for which the administrator has been authorized.

**restricted storage privilege.** An administrative privilege class that enables an administrator to control the allocation and use of storage resources only for the storage pools for which the administrator has been authorized.

**retention.** The amount of time, in days, that inactive backed up or archived files will be retained in the storage pool before they are deleted. The following copy group attributes define retention: retain extra versions, retain only version, retain version.

**retention period.** On an MVS server, a device class attribute that specifies how long files are retained on sequential access media. When used, ADSM passes this information to the MVS operating system to ensure that other tape management systems do not overwrite tape volumes that contain retained data.

**retrieve.** A function that allows users to copy an archive copy from the storage pool to the workstation. The archive copy in the storage pool is not affected. Contrast with *archive*.

**RLIO.** Record Level Input/Output.

**rollback.** To remove changes that were made to database files since the last commit point.

**root.** In the AIX and UNIX environments, the user name for the system user with the most authority.

**root user.** In the AIX and UNIX environments, an expert user who can log in and execute restricted commands, shut down the system, and edit or delete protected files. Also called the *superuser*.

**routing choice.** A choice in a pull-down menu that, when selected, brings up another window. See also *action choice*.

## S

**SAA.** Systems Application Architecture.

**schedule.** A database record that describes scheduled client operations or administrative commands. See *administrative command schedule* and *client schedule*.

**scheduling mode.** The type of scheduling operation set for the server and client. ADSM supports two scheduling modes for client operations: client-polling and server-prompted.

**SCSI.** Small computer system interface.

**selective backup.** A function that allows users to back up files or directories from a client domain 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 backup or archive processing. The value of this attribute determines whether processing continues, is retried, or is stopped. See *static*, *dynamic*, *shared static*, and *shared dynamic*.

**server.** A program that provides services to other programs (clients).

**server migration.** The process of moving data from one storage pool to the next storage pool as controlled by the high and low migration thresholds. See *high migration threshold* and *low migration threshold*.

**server options file.** A file that specifies processing options for communication methods, tape handling, pool sizes, language, and date, time, and number formats.

**server program.** The program that provides backup, archive, space management, and administrative services to clients. The server program must be at the necessary level to provide all of these services.

**server-prompted scheduling mode.** A client/server communication technique where the server contacts the client when work needs to be done.

**session resource usage.** The amount of wait time, CPU time, and space used or retrieved during a client session.

**shared dynamic.** A copy group serialization value that specifies that a file must not be modified during a backup or archive operation. ADSM attempts to retry the backup or archive operation a number of times; if the file is in use during each attempt, ADSM will back up or archive the file on its last try even though the file is in use. See also *serialization*. Contrast with *dynamic*, *shared static*, and *static*.

**shared static.** A copy group serialization value that specifies that the file must not be modified during backup or archive. ADSM will retry the backup or archive operation a number of times; if the file is in use during each attempt, ADSM will not back up or archive the file. See also *serialization*. Contrast with *dynamic*, *shared dynamic*, and *static*.

**shell.** In the AIX and UNIX environments, a software interface between a user and the operating system of a computer. Shell programs interpret commands and

user interactions on devices such as keyboards, pointing devices, and touch-sensitive screens and communicate them to the operating system.

**signal.** (1) A simple method of communication between two processes. One process can inform the other process when an event occurs. (2) In operating system operations, a method of inter-process communication that simulates software interrupts.

**signal handler.** A subroutine called when a signal occurs.

**SMIT.** System Management Interface Tool.

**SNA LU6.2.** Systems Network Architecture Logical Unit 6.2.

**socket.** (1) An endpoint for communication between processes or applications. (2) A pair consisting of TCP port and IP address, or UDP port and IP address.

**space 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 data storage. Administrators can be authorized with unrestricted or restricted storage privilege. See *restricted storage privilege* or *unrestricted storage privilege*.

**stub file.** A file that replaces the original file on a client node when the file is migrated from the client node to ADSM storage.

**superuser.** See *root user*.

**synchronized copy status.** Specifies that the volume is the only volume copy or is synchronized with other volume copies in the database or recovery log. When synchronized, mirroring has started.

**system privilege class.** An administrative privilege class that allows an administrator to issue all server commands.

**Systems Application Architecture (SAA).** Software interfaces, conventions, and protocols that provide a

framework for designing and developing applications that are consistent across systems.

**Systems Network Architecture (SNA).** A set of rules for data to be transmitted in a network. Application programs communicate with each other using a layer of SNA called advanced program-to-program communications (APPC).

## T

**tape.** A recording medium consisting of a long, narrow, flexible strip with a magnetic coating wound onto a reel or into a cartridge. See *cartridge* and *tape reel*.

**tape library.** (1) A term used to refer to a collection of tape cartridges. (2) An automated device that performs tape cartridge mounts and demounts without operator intervention.

**Tape Library Dataserver.** An automated tape library consisting of mechanical components, cartridge storage frames, IBM tape subsystems, and controlling hardware and software. The tape library dataserver performs tape cartridge mounts and demounts without operator intervention.

**tape reel.** A cylinder with flanges on which magnetic tape is wound. Devices such as the 3420 9-track tape device support tape reels. Contrast with *cartridge*.

**tape volume prefix.** A device class attribute that is the high-level-qualifier of the file name or the data set name in the standard tape label.

**task help.** A type of online help that provides a list of tasks that can be completed with a selected object. When you select a task, the help provides step-by-step information on how to complete the task.

**TCP/IP.** Transmission Control Protocol/Internet Protocol.

**Telnet.** In TCP/IP, the protocol that opens the connection to the system.

**Transmission Control Protocol/Internet Protocol (TCP/IP).** A set of communication protocols that support peer-to-peer connectivity functions for both local and wide area networks.

**trusted communication agent.** A program that performs communication tasks on behalf of the client or server, and ensures the security of the communications.

## U

**unit name.** On an MVS server, a device class attribute that specifies a group of tape devices used with the MVS server. A unit name can be a generic device type, an esoteric unit name, or a physical device.

**unrestricted policy privilege.** An administrative privilege class that enables an administrator to manage policy objects for any policy domain.

**unrestricted storage privilege.** An administrative privilege class that enables an administrator to control the database, recovery log, and all storage pools.

**utilization.** The percent of assigned capacity used by the database or recovery log at a specific point of time.

## V

**validate.** The process of ensuring that the active policy set contains a default management class and reporting on copy group definition errors.

**version.** The maximum number of backup copies retained for files and directories. The following copy group attributes define version criteria: versions data exists and versions data deleted.

**Virtual Machine (VM).** One of the family of IBM operating systems for the System/370 or System/390 processor, including VM/ESA, VM/XA, VM/SP, and VM/HPO. VM is one of the supported server environments.

**Virtual Storage Extended (VSE).** One of the family of IBM operating systems for the System/370 or System/390 processor, including VSE/ESA. VSE is one of the supported server environments.

**VM.** Virtual Machine.

**volume.** The basic unit of storage for the database, recovery log, or a storage pool. A volume can be an LVM logical volume, a standard file system file, a tape cartridge, or an optical cartridge. Each volume is identified by a unique volume identifier. See *database volume* and *storage pool volume*.

**volume history file.** A file that contains information about: volumes used for database backups and database dumps; volumes used for export of administrator, node, policy, or server data; and sequential access storage pool volumes that have been added, reused, or deleted. The information is a copy of the same types of volume information in the ADSM database.



**volume set.** An entire image of the database or recovery log, as displayed on the administrative graphical user interface.

**VSE.** Virtual Storage Extended.

## W

**WDSF/VM.** Workstation Data Save Facility/Virtual Machine.

**wildcard character.** A character or set of characters used to specify an unknown number or set of characters in a search string. Also called *pattern-matching character*.

**window.** A part of a display screen with visible boundaries in which information is displayed.

**windowed interface.** A type of user interface that is either a graphical user interface or a text based inter-

face. The text based interface maintains a close affinity to the graphical user interface, including action bars and their associated pull-downs and windows. See *graphical user interface*.

**workstation.** A personal computer system capable of maintaining data files.

**Workstation Data Save Facility/Virtual Machine (WDSF/VM).** The predecessor product to ADSTAR Distributed Storage Manager.

## 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

- absolute mode, description of 130
- access mode 214
- accounting record
  - description of 59
  - determining for storage pool 194, 197
  - monitoring 59
- ACTIVATE POLICYSET command 37, 134
- activity log
  - adjusting the size 57
  - description of 56
  - monitoring 56
  - querying 58
  - setting the retention period 57
- administrative client
  - description of 1
  - options 24
  - viewing information after IMPORT or EXPORT 245
- administrative command line interface 22
- administrative commands
  - AUDIT LICENSE 33, 92
  - AUDIT VOLUME 43, 222, 234
  - BACKUP DB 46, 273
  - BACKUP DEVCONFIG 46
  - BACKUP STGPOOL 42, 280
  - BACKUP VOLHISTORY 46
  - DEFINE DEVCLASS 41, 168
  - DEFINE DOMAIN 36
  - DEFINE SCHEDULE 149
  - DEFINE STGPOOL 42, 196, 197
  - DEFINE VOLUME 216
  - DELETE DEVCLASS 41
  - DELETE LOGVOLUME 82
  - DELETE VOLHISTORY 270
  - DELETE VOLUME 229, 230
  - DISMOUNT VOLUME 235
  - EXTEND LOG 275
  - GRANT AUTHORITY 94
  - HELP 61
  - MOVE DATA 225, 234
  - QUERY ACTLOG 58
  - QUERY CONTENT 234
  - QUERY DB 80
  - QUERY DBBACKUPTRIGGER 274
  - QUERY DEVCLASS 41
  - QUERY LICENSE 92
  - QUERY MOUNT 235
  - QUERY OCCUPANCY 208, 209
  - QUERY OPTION 56
  - QUERY PROCESS 228
  - QUERY STGPOOL 42, 202, 204, 254
- administrative commands (*continued*)
  - QUERY VOLHISTORY 270
  - QUERY VOLUME 217, 218, 228, 234
  - REGISTER ADMIN 93
  - REGISTER LICENSE 92
  - RENAME ADMIN 100
  - RESET DBMAXUTILIZATION 70
  - RESET LOGMAXUTILIZATION 70
  - RESTORE STGPOOL 47, 280, 282
  - RESTORE VOLUME 43, 284
  - SET ACCOUNTING 58
  - SET ACTLOGRETENTION 57
  - SET AUTHENTICATION 93
  - SET LICENSEAUDITPERIOD 92
  - SET LOGMODE 47
  - SET PASSEXP 93
  - SET SCHEDMODES 144
  - SET SERVERNAME 56
  - UPDATE ADMIN 93
  - UPDATE DBBACKUPTRIGGER 274
  - UPDATE DEVCLASS 41
  - UPDATE SCHEDULE 149
  - UPDATE VOLUME 216
- administrative privilege class
  - analyst 98
  - description of 94
  - granting authority 94
  - operator 97
  - policy 95, 96
  - revoking all 98
  - storage 96, 97
  - system 95
- administrator
  - authorizing to manage a policy domain 94
  - description of 1
  - locking 101
  - managing registration 89
  - querying 102
  - registering 93
  - removing 101
  - renaming 100
  - unlocking 101
  - updating 93, 94
  - viewing information about 102
- analyst privilege class
  - description of 98
  - granting 98
  - revoking 98
- API
  - compression option 110
  - deletion option 111
  - description of 1, 7, 26

- application client, registering 110
- application programming interface
  - description of 1, 26
  - information about xvii
  - registering to server 110
  - using 7
- archive
  - amount of space used 209
  - defining criteria 124
  - description of 4
  - processing 121
- archive copy group
  - defining 132, 133
  - deleting 138
  - description of 116
- archive file management 114
- archiving a file 114, 121
- ARCTVEXT exit 238
- ASSIGN DEFMGMTCLASS command 37, 134
- assigned capacity 69, 77
- association
  - defining 153
  - deleting 154
  - description of 148
  - querying 154
  - viewing information about 154
- audit license 95
- AUDIT LICENSE command 92
- AUDIT VOLUME command 220, 222, 234
- authentication, client/server 92
- authority
  - granting to administrators 94
  - revoking 98

## B

- backup
  - a file 114, 120, 121
  - amount of space used 209
  - database 273, 274
  - defining criteria 124
  - description of 4
  - file management 114
  - incremental 114, 120
  - selective 114, 121
  - static 129, 133
  - storage pool 279
  - when to perform 267
- backup copy group
  - defining 128, 132
  - deleting 138
  - description of 116
  - frequency 120
  - mode 120
  - serialization 120

- BACKUP DB command 46, 273
- BACKUP DEVCONFIG command 46, 272
- backup period, specifying for incremental 145
- BACKUP STGPOOL command 42, 280
- BACKUP VOLHISTORY command 46, 270
- backup-archive client
  - description of 1
  - registering 104
- batch mode 23
- binding
  - description of 118
  - file to a management class 118
- buffer pool 83
- BUFPOOLSIZ option 84

## C

- cache
  - description of 4
  - enabling for disk storage pools 185, 195
  - monitoring utilization on disk 207
- CANCEL PROCESS command 29, 42, 54, 206
- CANCEL SESSION command 29, 52
- capacity, assigned 69, 77
- CARTRIDGE device class, defined 163
- cartridge tape, setting the retention period for files 166
- central scheduling
  - controlling the workload 145
  - coordinating 145
  - description of 4, 141
- class, administrator privilege
  - analyst 98
  - description of 94
  - granting authority 94
  - operator 97
  - policy 95, 96
  - revoking all 98
  - storage 96, 97
  - system 95
- class, device
  - 4MM 161, 163
  - 8MM 161, 163
  - amount of space used 209
  - CARTRIDGE 163
  - defining for backup 269
  - deleting 169
  - description of 4
  - DISK 161
  - FILE 161
  - LAN3995 161
  - MO35 161
  - optical 161
  - querying 168
  - selecting for import and export 243
- class, policy privilege
  - description of restricted 96

- class, policy privilege (*continued*)
  - description of unrestricted 95
  - granting restricted 96
  - granting unrestricted 95
  - reducing unrestricted to restricted 99
  - revoking 98
- class, storage privilege
  - description of restricted 97
  - description of unrestricted 96
  - granting restricted 97
  - granting unrestricted 96
  - reducing unrestricted to restricted 100
  - revoking 98
- client
  - administrative 1
  - application 26
  - backup-archive 4
  - space-management 1
- client migration 121, 122
- client node
  - amount of space used 208
  - locking 106
  - managing registration 89, 102
  - querying 107
  - registering 104, 110
  - removing 110
  - renaming 106
  - setting password authentication 93
  - setting scheduling mode 144
  - unlocking 106
  - updating 105
  - viewing information about 107
- client queries to the server, setting the frequency 147
- client session
  - canceling 52
  - managing 51
  - querying 51
  - viewing information about 51
- client system options file 111
- client-polling scheduling 142, 147
- client/server, description of 1
- closed registration
  - description of 104
  - setting 103
- collocation
  - definition 186, 195, 198
  - description of 4
  - determining whether to use collocation 186, 195, 198
  - enabling for sequential storage pool 186, 195, 198
  - how it affects reclamation 191
  - how the server selects volumes when disabled 189
  - how the server selects volumes when enabled 188
  - migration thresholds 184
- command retry attempts
  - setting the amount of time between 148
- command retry attempts (*continued*)
  - setting the number of 147
- command-line interface, using 7
- commands, administrative
  - AUDIT LICENSE 33, 92
  - AUDIT VOLUME 43, 222, 234
  - BACKUP DB 46, 273
  - BACKUP DEVCONFIG 46
  - BACKUP STGPOOL 42, 280
  - BACKUP VOLHISTORY 46
  - DEFINE DEVCLASS 41, 168
  - DEFINE DOMAIN 36
  - DEFINE SCHEDULE 149
  - DEFINE STGPOOL 42, 196, 197
  - DEFINE VOLUME 216
  - DELETE DEVCLASS 41
  - DELETE LOGVOLUME 82
  - DELETE VOLHISTORY 270
  - DELETE VOLUME 229, 230
  - DISMOUNT VOLUME 235
  - EXTEND LOG 275
  - GRANT AUTHORITY 94
  - HELP 61
  - MOVE DATA 225, 234
  - QUERY ACTLOG 58
  - QUERY CONTENT 234
  - QUERY DB 80
  - QUERY DBBACKUPTRIGGER 274
  - QUERY DEVCLASS 41
  - QUERY LICENSE 92
  - QUERY MOUNT 235
  - QUERY OCCUPANCY 208, 209
  - QUERY OPTION 56
  - QUERY PROCESS 228
  - QUERY STGPOOL 42, 202, 204, 254
  - QUERY VOLHISTORY 270
  - QUERY VOLUME 217, 218, 228, 234
  - REGISTER ADMIN 93
  - REGISTER LICENSE 92
  - RENAME ADMIN 100
  - RESET DBMAXUTILIZATION 70
  - RESET LOGMAXUTILIZATION 70
  - RESTORE STGPOOL 47, 280, 282
  - RESTORE VOLUME 43, 284
  - SET ACCOUNTING 58
  - SET ACTLOGRETENTION 57
  - SET AUTHENTICATION 93
  - SET LICENSEAUDITPERIOD 92
  - SET LOGMODE 47
  - SET PASSEXP 93
  - SET SCHEDMODES 144
  - SET SERVERNAME 56
  - UPDATE ADMIN 93
  - UPDATE DBBACKUPTRIGGER 274
  - UPDATE DEVCLASS 41
  - UPDATE SCHEDULE 149

commands, administrative (*continued*)

- UPDATE VOLUME 216
- COMMIT command 28
- compression 103, 104, 110, 165
- compression option, API 110
- configuration, device
  - files
    - backup 269, 271
    - dsmserv.opt 50, 269, 271
  - information 269, 271
- console mode 22, 245
- contextual help menu
  - accessing 18
  - description of 15
- COPY DOMAIN command 36, 126
- copy group
  - deleting 138
  - description of 4
- COPY MGMTCLASS command 36, 128
- COPY POLICYSET command 36, 127
- COPY SCHEDULE command 39, 152

## D

data

- considering user needs for recovering 160
- exporting 241
- importing 241
- data compression 165
- data movement, querying about the process 228
- data storage
  - considering user needs for recovering 160
  - deleting files from 228
  - evaluating 159
  - managing 157
  - monitoring 219
  - tailoring definitions 256
  - understanding how files are stored 157

database

- adding space to 75
- allocating space 65
  - procedure 71
- available space 68
- backup 273
- backup trigger 273
- buffer pool 83, 84
- committing data to 83
- defining a volume 75
- defining mirrored volumes 263
- deleting a volume 82
- deleting space 78
- description of 3, 63
- determining how much space is allocated 67
- ensuring integrity of 5
- estimating the amount of space needed 66
- logical volume 67

database (*continued*)

- managing 63
- mirroring 262
- monitoring space 70
- monitoring the buffer 84
- optimizing performance 83
- querying the buffer pool 84
- recovering 276
- reducing capacity 82
- references to storage pools 65
- resetting buffer pool statistics 85
- restoring 275
- transactions 63, 64, 65
- viewing information about 84
- volume placement 71
- database backup
  - when to perform 267
- database backup and recovery
  - defining device classes 269
  - full backup 267
  - incremental backup 267
  - point-in-time 265, 266
  - roll-forward 266, 267
  - to most current state 266
  - trigger 268
- database backup trigger and roll-forward mode 266
- database recovery
  - example recovery procedures 285
  - general strategy 261
  - methods 261
  - offline utilities 262
  - providing 261
  - salvage utilities 284
  - when to backup 267
- day of the week, description of 149
- default management class, description of 115
- default management classes 117
- DEFINE ASSOCIATION command 40, 153
- DEFINE COPYGROUP command 36, 128, 132, 133
- DEFINE DBBACKUPTRIGGER 47, 269, 273
- DEFINE DBCOPY command 46, 263
- DEFINE DBVOLUME command 31, 75
- DEFINE DEVCLASS command 167
- DEFINE DOMAIN command 36
- DEFINE LOGCOPY command 46, 263
- DEFINE LOGVOLUME command 31, 75
- DEFINE MGMTCLASS command 36, 127
- DEFINE POLICYSET command 36, 127
- DEFINE SCHEDULE command 39, 149
- DEFINE STGPOOL command 196, 197
- DEFINE VOLUME command 43, 216
- delete
  - empty volume 229
  - files 228
  - storage volumes 229
  - volume with residual data 230

- DELETE ASSOCIATION command 40, 154
- DELETE COPYGROUP command 37, 138
- DELETE DBBACKUPTRIGGER 47, 273
- DELETE DBVOLUME command 32, 82
- DELETE DEVCLASS command 169
- DELETE DOMAIN command 38, 139
- DELETE EVENT command 40, 156
- DELETE FILESPACE command 34, 109
- DELETE LOGVOLUME command 32, 82
- DELETE MGMTCLASS command 37, 139
- DELETE POLICYSET command 37, 139
- DELETE SCHEDULE command 40, 153
- DELETE STGPOOL command 42, 210
- DELETE VOLHISTORY command 46, 270
- DELETE VOLUME command 43, 229, 230
- deletion exit
  - parameter list 238
  - registers on entry to 238
- DELETIONEXIT option 237, 238
- DEVCONFIG option 271
- device class
  - 4MM 161, 163
  - 8MM 161, 163
  - amount of space used 209
  - CARTRIDGE 163
  - defining for backup 269
  - deleting 169
  - description of 4
  - DISK 161
  - FILE 161
  - LAN3995 161
  - MO35 161
  - optical 161
  - querying 168
  - selecting for import and export 243
- device configuration
  - files
    - backup 269, 271
    - dsmserv.opt 50, 269, 271
    - information 269, 271
- device configuration files, recreating 272
- device type, selecting 161
- DFRMM 236
- DFSMShsm, ARCTVEXT Exit 238
- DFSMSrmm 236
- DISABLE command 29, 53
- disaster recovery
  - example recovery procedures 285
  - general strategy 261
  - methods 261
  - offline utilities 262
  - providing 261
  - salvage utilities 284
  - when to backup 267
- DISK device class, defined 161
- disk storage pool
  - estimating space 192
  - estimating space for archived files 193
  - estimating space for backed up files 193
- DISMOUNT DEVICE command 235
- DISMOUNT VOLUME command 44, 235
- dismounting
  - tape drives 235
  - tape volumes 235
- documentation, user xv
- DSMSERV AUDITDB utility 48, 284
- DSMSERV DUMPDB utility 47, 284
- DSMSERV INSTALL utility 47, 284
- DSMSERV LOADDDB utility 47, 284
- DSMSERV RESTORE DB command 47, 276, 278
- dsmserv.opt 50, 269, 271
- dynamic serialization, description of 129, 133

## E

- edit, description of menu choices 10
- ENABLE command 29, 53
- error log, monitoring 58
- error states for tape volumes 234
- esoteric unit name, specifying 165
- estimated capacity for tape volumes, specifying 166
- event record
  - deleting 156
  - description of 141
  - removing from the database 156
- event retention period 156
- event, description of 141
- exits, installation-wide tape, deletion 237
- EXPINTERVAL option 119, 135
- expiration date
  - setting for files on tape 165
  - setting schedule 149
- expiration processing 119, 135
- EXPIRE INVENTORY command 119, 135
- export
  - monitoring 244
  - planning for sequential media 243
  - PREVIEW parameter 242
  - querying about a process 244
  - querying the activity log 246
  - using scratch media 243
  - viewing information about a process 244
- EXPORT ADMIN command 45, 249
- EXPORT command 245
- EXPORT NODE command 45, 250
- EXPORT POLICY command 45, 250
- EXPORT SERVER command 45, 243, 249
- exporting
  - administrator data 249
  - client node data 250
  - data to tape 247

exporting (*continued*)  
description of 5, 241  
policy data 250  
server data 249  
EXTEND DB command 31, 77  
EXTEND LOG command 31, 77

## F

file data, importing 241  
FILE device class, defined 161  
file expiration date 165  
file retention period 166  
file size, determining maximum for storage pool 195  
filesystem  
deleting 109  
description of 4, 107  
querying 108  
viewing information about 108  
frequency, description of 129

## G

general help menu  
accessing 16  
description of 15  
GRANT AUTHORITY command 33  
graphical user interface  
accessing contextual help 18  
accessing general help 16  
accessing help from windows 16  
accessing task help 17  
administration window 9  
completing tasks when using 11  
ending a client session 14  
starting from a window 9  
starting from an icon 9  
using 7  
using hypertext links 19  
using menu choices 10  
using online help 15  
using the search facility 20

## H

HALT command 28, 50  
halting the server 50  
help  
description of 16  
description of menu choices 10  
HELP command 30  
help index menu, description of 15  
hierarchial storage management  
automatic  
demand 114  
threshold 114

hierarchial storage management (*continued*)  
premigration 114  
reconciliation 115  
selective 114  
stub file 114  
HSM  
See hierarchial storage management

## I

I/O error, tape 234  
import  
monitoring 244  
PREVIEW parameter 242, 252  
querying about a process 244  
querying the activity log 246  
recovering from an error 259  
viewing information about a process 244  
IMPORT ADMIN command 45, 251  
IMPORT command 245  
IMPORT NODE command 45, 251, 258  
IMPORT POLICY command 45, 251  
IMPORT SERVER command 45, 251, 258  
importing  
data 251  
data storage definitions 254, 256  
date of creation 257  
description of 5, 241  
directing messages to an output file 255  
duplicate file spaces 257  
file data 257  
policy definitions 254  
server control data 256  
subsets of information 259  
include-exclude file  
description of 4  
for policy environment 124  
incremental backup  
file eligibility for 120  
incremental backup frequency, specifying 145  
initial start date, description of 149  
initial start time, description of 149  
interactive mode 23  
interface, application programming  
description of 1, 26  
information about xvii  
registering to server 110  
using 7  
interface, command-line 7  
interface, graphical user  
accessing contextual help 18  
accessing general help 16  
accessing help from windows 16  
accessing task help 17  
administration window 9  
completing tasks when using 11



interface, graphical user (*continued*)  
  ending a client session 14  
  starting from a window 9  
  starting from an icon 9  
  using 7  
  using hypertext links 19  
  using menu choices 10  
  using online help 15  
  using the search facility 20

## L

license  
  compliance 92  
  example 92  
  features  
    for additional clients 90  
    for device module support 91  
    for other clients 90  
  monitoring 92  
  using 89  
LOCK ADMIN command 34, 101  
LOCK NODE command 34, 106  
log mode  
  normal 273, 275  
  roll-forward 273, 275  
  setting 275  
LOGPOOLSIZE option 84

## M

MACRO command 24  
macros 24  
management class  
  assigning a default 134  
  associating a file with 118  
  binding a file to 118  
  controlling user access 116  
  copying 123, 127  
  default 117  
  defining 127  
  deleting 139  
  description of 4, 115, 116  
  rebinding a file 119  
  updating 123, 128  
management class configuration 117  
maximum extension 76  
media loss, recovery from 284  
menu bar choices 10  
menu choices 10  
migrating a file 114, 121  
migration  
  canceling 206  
  defining threshold for disk storage pool 183  
  defining threshold for tape storage pool 184  
  description of 4

migration (*continued*)  
  description of when files are migrated 181  
  monitoring thresholds 204  
  providing users with immediate access to files on disk 183  
  threshold  
    high 181  
    low 181  
  updating the storage pool to complete the process 206  
mirrored volume  
  description of 264  
  querying 264  
  viewing information about 264  
mirroring  
  advantages 262  
  database 262  
  defining volumes 263  
  description of 5  
  recovery log 262, 263  
  recovery procedure 265  
mode  
  description of 129  
  scheduling 144  
modified mode, description of 129  
mount  
  limit 163, 168  
  messages, tape 233  
  mode 23  
  query 235  
  retention period 164  
  wait period 164  
MOVE DATA command 43, 225, 234  
MVS esoteric unit name, specifying 165

## O

occupancy, querying 208  
online help  
  accessing from windows 16  
  finding information 19  
  using 15  
  using hypertext links 19  
  using the search facility 20  
open registration  
  description of 103  
  setting 103  
operator privilege class  
  description of 97  
  granting 97  
  revoking 98  
optical device class, defined 161  
option  
  DEVCONFIG 271  
  VOLUMEHISTORY 270

- options file, server
  - See dsmserv.opt
- options, querying
  - BUFPOOLSIZE 84
  - DELETIONEXIT 237, 238
  - LOGPOOLSIZE 84
  - VIRTUALMOUNTPOINT 107

## P

- page, description of 83
- password
  - setting authentication for a client 93
  - setting expiration 93
- period, specifying for an incremental backup 145
- policies, managing AD5M 113
- policy definitions, importing 254
- policy domain
  - creating 126
  - deleting 139
  - description of 115, 116
  - updating 123, 126
- policy objects
  - deleting 138
  - description of 115
  - querying 136
- policy operations 113
- policy privilege class
  - description of restricted 96
  - description of unrestricted 95
  - granting restricted 96
  - granting unrestricted 95
  - reducing unrestricted to restricted 99
  - revoking 98
- policy set
  - activating 134
  - copying 123, 127
  - defining 127
  - deleting 139
  - description of 115, 116
  - updating 127
  - validating 134, 135
- pool, storage
  - amount of space used 209
  - backup and recovery 279
  - copy 172, 173
  - creating a hierarchy 178
  - defining 194
  - defining for disk 196
  - defining for tape 196
  - deleting 210
  - description of 172
  - determining access mode 194, 197
  - determining maximum file size 195
  - determining whether to use collocation 186, 195, 198

- pool, storage (*continued*)
  - enabling cache for disk 185, 195
  - estimating space for archived files on disk 193
  - estimating space for backed up files on disk 193
  - estimating space for disk 192
  - estimating space for sequential 194
  - estimating space in multiple 178
  - managing 171
  - monitoring 202
  - moving files 225
  - moving files between 226
  - overview 158
  - primary 172
  - querying 202
  - random access 172
  - references in recovery log 65
  - restore 280, 282
  - sequential access 172
  - updating 194
  - updating for disk 197
  - using cache on disk 185, 195
  - viewing information about 202
- premigration 115
- PREVIEW parameter 242, 252
- priority 149
- privilege class, administrator
  - analyst 98
  - description of 94
  - granting authority 94
  - operator 97
  - policy 95, 96
  - revoking all 98
  - storage 96, 97
  - system 95
- privilege class, policy
  - description of restricted 96
  - description of unrestricted 95
  - granting restricted 96
  - granting unrestricted 95
  - reducing unrestricted to restricted 99
  - revoking 98
- process canceling 54
- processing tapes 233
- product information menu, description of 15
- programming interface notice xi
- protection status, determining for cartridge tape 166
- publications xv

## Q

- QUERY ACTLOG command 30, 58, 246
- QUERY ADMIN command 34, 102
- QUERY ASSOCIATION command 40, 154
- QUERY CONTENT command 43, 223, 234
- QUERY COPYGROUP command 37, 136, 256

QUERY DB command 32, 80, 84  
 QUERY DBBACKUPTRIGGER command 47, 274  
 QUERY DBVOLUME command 31, 46, 79, 264  
 QUERY DEVCLASS command 168, 243  
 QUERY DOMAIN command 37, 138  
 QUERY EVENT command 40, 155  
 QUERY FILESPACE command 34, 108  
 QUERY LICENSE command 33, 92  
 QUERY LOG command 32, 86  
 QUERY LOGVOLUME command 31, 46, 79, 264  
 QUERY MGMTCLASS command 37, 137  
 QUERY MOUNT command 44, 235  
 QUERY NODE command 34, 107  
 QUERY OCCUPANCY 208, 209  
 QUERY OCCUPANCY command 42, 208  
 QUERY OPTION command 30  
 QUERY POLICYSET command 37, 137  
 QUERY PROCESS command 29, 43, 45, 54, 206, 228, 244  
 QUERY SCHEDULE command 40, 153  
 QUERY SESSION command 29, 51  
 QUERY STATUS 30, 55  
 QUERY STGPOOL 42, 202, 204, 207  
 QUERY VOLHISTORY 46  
 QUERY VOLUME command 43, 217, 218, 228, 234  
 querying for general information 218  
 querying policy objects 136  
 querying storage volumes 218  
 QUIT command 23

## R

randomize, description of 146  
 read-only access mode 214  
 read/write access mode 214  
 rebinding  
   description of 119  
   file to a management class 119  
 recalling a file  
   selective 115  
   transparent 115  
 reclamation  
   affect of collocation on 191  
   description of 4  
   setting a threshold for sequential storage pool 190, 195, 198  
   threshold 4  
 reclamation threshold, setting for sequential storage pool 190, 195, 198  
 recovering storage pools 279  
 recovering the database 276  
 recovery from disaster  
   See disaster recovery  
 recovery log  
   adding space to 75  
   allocating space 65  
   procedure 71  
   recovery log (*continued*)  
     available space 68  
     buffer pool 86  
     consistent database image 63  
     defining a volume 75  
     defining mirrored volumes 263  
     deleting a volume 82  
     deleting space 78  
     description of 3, 63  
     determining how much space is allocated 67  
     estimating the amount of space needed 66  
     logical volume 67  
     managing 63  
     mirroring 262, 263  
     monitoring space 67  
     monitoring the buffer pool 86  
     optimizing performance 83  
     querying the buffer pool 86  
     reducing capacity 82  
     references to storage pools 65  
     size of 275  
     viewing information about 86  
     volume placement 71  
     when to backup 267  
   recovery log mode  
     normal 273, 275  
     roll-forward 273, 275  
     setting 275  
   recovery, disaster  
     example recovery procedures 285  
     general strategy 261  
     methods 261  
     offline utilities 262  
     providing 261  
     salvage utilities 284  
     when to backup 267  
 REDUCE DB command 31, 82  
 REDUCE LOG command 32, 82  
 REGISTER ADMIN command 33, 93  
 REGISTER LICENSE command 33, 92  
 REGISTER NODE command 34, 35, 104  
 registering a workstation 110  
 registration  
   closed 104  
   description of 103  
   managing client node 102  
   managing for a client node 89  
   managing for an administrator 89  
   open 103  
   setting for a client node 103  
 REMOVE ADMIN command 34, 101  
 REMOVE NODE command 35, 110  
 RENAME ADMIN command 33, 100  
 RENAME FILESPACE command 45  
 RENAME NODE command 34, 106

- renaming an administrator ID 100
- RESET BUFPOOL command 32, 85
- RESET DBMAXUTILIZATION command 31, 70
- RESET LOGMAXUTILIZATION command 31, 70
- restarting the server 50
- RESTORE STGPOOL command 42, 280, 282
- RESTORE VOLUME command 43, 284
- restoring a file 114
- restoring the database
  - point-in-time 276
  - to its most current state 278
- restricted policy privilege
  - granting 96
  - reducing 100
- restricted storage privilege
  - granting 97
  - reducing 100
- retain extra versions, description of 131
- retain only version, description of 132
- retention grace period
  - description of archive 126
  - description of backup 126
  - using archive 126
  - using backup 126
- retrieving a file 114
- REVOKE AUTHORITY command 33, 98
- roll-forward recovery
  - database backup trigger 266
  - mirroring recovery log 266
  - recovery log 266
- ROLLBACK command 28
- root user tasks
  - define compression 103
  - registering your workstation 110

## S

- schedule
  - associating client node 153
  - client options to use 151
  - coordinating 145
  - copying 152
  - day of the week 149
  - defining 149
  - deleting 153
  - description of 141
  - expiration date 149
  - files to process 151
  - frequency of service 149
  - initial start date 149
  - initial time 149
  - managing associations 148
  - priority 149
  - querying 153
  - randomize 146
  - startup window 146, 149
- schedule (*continued*)
  - type of action 151
  - updating 149
  - viewing information about 153
- schedule event
  - managing 155
  - querying 155
  - viewing information about 155
- scheduled operations, setting the maximum 145
- scheduler workload, controlling 145
- scheduling mode
  - client-polling 142
  - description of 141
  - selecting 144
  - server-prompted 142
  - setting on a client node 144
  - setting on the server 144
- scheduling, central
  - controlling the workload 145
  - coordinating 145
  - description of 4, 141
- scratch volumes, determining use 178, 195, 198
- search facility, using 20
- selected, description of menu choices 10
- selective backup 114, 121
- selective recall 115
- sequential storage pool, estimating space 194
- server
  - canceling process 54
  - console session 25
  - description of 1
  - description of how tapes are selected 233
  - disabling 53
  - disabling access 53
  - enabling 53
  - enabling access 53
  - halting 50
  - managing operation 49
  - managing processes 53
  - querying about processes 54
  - querying options 56
  - querying status 55
  - restarting 50
  - scheduling mode 39
  - setting the name 56
  - starting 49
  - stopping 50
  - viewing information about 55
  - viewing information about processes 54
- server console, description of 94
- server options file
  - See dsmserv.opt
- server-prompted scheduling 142
- session
  - canceling 52
  - setting the maximum percentage for scheduled operations 145

- SET ACCOUNTING command 30, 59
- SET ACTLOGRETENTION command 30, 57
- SET AUTHENTICATION command 33, 93
- SET EVENTRETENTION command 40, 156
- SET LICENSEAUDITPERIOD command 33
- SET LOGMODE command 47, 274
- SET MAXCMDRETRIES command 39, 147
- SET MAXSCHEDSESSIONS command 145
- SET PASSEXP command 33, 93
- SET QUERYSCHEDPERIOD command 39, 147
- SET RANDOMIZE command 39, 146
- SET REGISTRATION command 34, 103
- SET RETRYPERIOD command 39, 148
- SET SCHEDMODES command 144
- SET SERVERNAME command 30
- setting a password 110
- setting compression 110
- shared dynamic serialization, description of 129, 133
- shared static serialization, description of 129, 133
- SMF records 59
- space
  - planning how much to allocate 71
  - procedure for allocating 71
  - unusable 68
  - usable 68
- space management
  - See hierarchial storage management
- standard management class, copying 128
- standard storage management policies, using 122
- starting the server 49
- startup window, description of 146
- stopping the server 50
- storage hierarchy
  - defining in reverse order 197
  - establishing 178
- storage management policies
  - description of 4, 115
  - managing 113
  - tailoring 123
  - using standard 122
- storage occupancy, querying 208
- storage pool
  - amount of space used 209
  - backup and recovery 279
  - copy 172, 173
  - creating a hierarchy 178
  - defining 194
  - defining for disk 196
  - defining for tape 196
  - deleting 210
  - description of 172
  - determining access mode 194, 197
  - determining maximum file size 195
  - determining whether to use collocation 186, 195, 198
  - enabling cache for disk 185, 195

- storage pool (*continued*)
  - estimating space for archived files on disk 193
  - estimating space for backed up files on disk 193
  - estimating space for disk 192
  - estimating space for sequential 194
  - estimating space in multiple 178
  - managing 171
  - monitoring 202
  - moving files 225
  - moving files between 226
  - overview 158
  - primary 172
  - querying 202
  - random access 172
  - references in recovery log 65
  - restore 280, 282
  - sequential access 172
  - updating 194
  - updating for disk 197
  - using cache on disk 185, 195
  - viewing information about 202
- storage pool backup
  - full 166, 280
  - incremental 280
- storage privilege class
  - description of restricted 97
  - description of unrestricted 96
  - granting restricted 97
  - granting unrestricted 96
  - reducing unrestricted to restricted 100
  - revoking 98
- storage volume
  - auditing 220
  - contents 223
  - information about 218
  - monitoring use 217
  - overview 158
- stub file 114
- system privilege class
  - description of 95
  - granting 95
  - reducing 99
  - revoking 98

## T

- tape
  - controlling mount messages 233
  - deletion exit 237
  - determining protection status for CARTRIDGE 166
  - dismounting 235
  - error states 234
  - exporting data 247
  - handling mount requests 235
  - mount request 234
  - operations, managing 233

- tape (*continued*)
  - planning for exporting data 243
  - processing 233
  - querying mounted 235
  - requesting 234
  - scratch, determining use 178, 195, 198
  - setting a mount limit 163
  - setting mount retention period 164
  - specifying estimated capacity 166
- tape exits, installation-wide, deletion 237
- tape management system
  - defining ADSM to 235
  - DFRMM 235
  - DFSMSrmm 235
  - setting up 235
- task help menu
  - accessing 17
  - description of 15
- TRACE commands 58
- trademarks xiii
- transactions, database 63, 64, 65
- transparent recall 115

## U

- unavailable access mode 214
- unavailable error state, tape 234
- unavailable tape volumes, recovering 234
- unit name, specifying 165
- UNLOCK ADMIN command 34, 101
- UNLOCK NODE command 34, 106
- unplanned shutdown 50
- unrestricted policy privilege
  - granting 95
  - reducing to restricted 99
- unrestricted storage privilege
  - granting 96
  - reducing to restricted 100
- unusable space 68
- UPDATE ADMIN command 33, 94
- UPDATE COPYGROUP command 36, 128, 132
- UPDATE DBBACKUPTRIGGER command 47
- UPDATE DOMAIN command 36, 126
- UPDATE MGMTCLASS command 36, 128
- UPDATE NODE command 34, 105
- UPDATE POLICYSET command 36, 127
- UPDATE SCHEDULE command 39, 149
- UPDATE STGPOOL command 42
- UPDATE VOLUME command 43, 216
- usable space 68
- user documentation xv
- using help menu, description of 15
- utilization
  - description of 69
  - monitoring 70

## V

- VALIDATE POLICYSET command 37, 134
- VARY command 29, 55
- varying volumes on or off line 55
- versions data deleted, description of 131
- versions data exists, description of 130
- view, description of menu choices 10
- VIRTUALMOUNTPOINT option 107
- volume
  - auditing 220
  - auditing considerations 221
  - defining for database 75
  - defining for recovery log 75
  - defining to storage pools 216
  - deleting 229, 230
  - detailed report 224
  - disk storage 216
  - disk storage pool 222
  - error states on tape 234
  - monitoring movement of data 228
  - monitoring use 217
  - moving files between 225
  - preparing for data storage 213
  - querying 218
  - querying contents 223
  - querying for general information 218
  - random access storage pools 172, 178, 215
  - recovery using mirroring 265
  - sequential 216
  - setting access mode 214
  - standard report 224
  - updating to storage pools 216
  - varying 55
- volume copy
  - allocating to separate disks 263
  - description of 262
- volume history backup
  - and point-in-time 268
- volume history files, establishing 270
- VOLUMEHISTORY option 268, 270

## W

- Windows
  - gui administration 9
  - using menu choices 10
- workstation, registering 110



---

# Communicating Your Comments to IBM

ADSTAR Distributed Storage Manager  
for MVS  
Administrator's Guide  
Version 2

Publication No. SH26-4039-00

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](mailto: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 MVS  
Administrator's Guide  
Version 2  
Publication No. SH26-4039-00**

**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.

\_\_\_\_\_




Cut or Fold  
Along Line


Fold and Tape

Please do not staple

Fold and Tape



NO POSTAGE  
NECESSARY  
IF MAILED IN THE  
UNITED STATES



# BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

International Business Machines Corporation  
Information Development  
Department 61C  
9000 South Rita Road  
TUCSON AZ 85775-4401



Fold and Tape

Please do not staple

Fold and Tape

Cut or Fold  
Along Line





Program Number: 5655-119



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

SH26-4039-00



*Spine information:*



**ADSTAR Distributed Storage Manager  
for MVS**

**Administrator's Guide**

*Version 2*