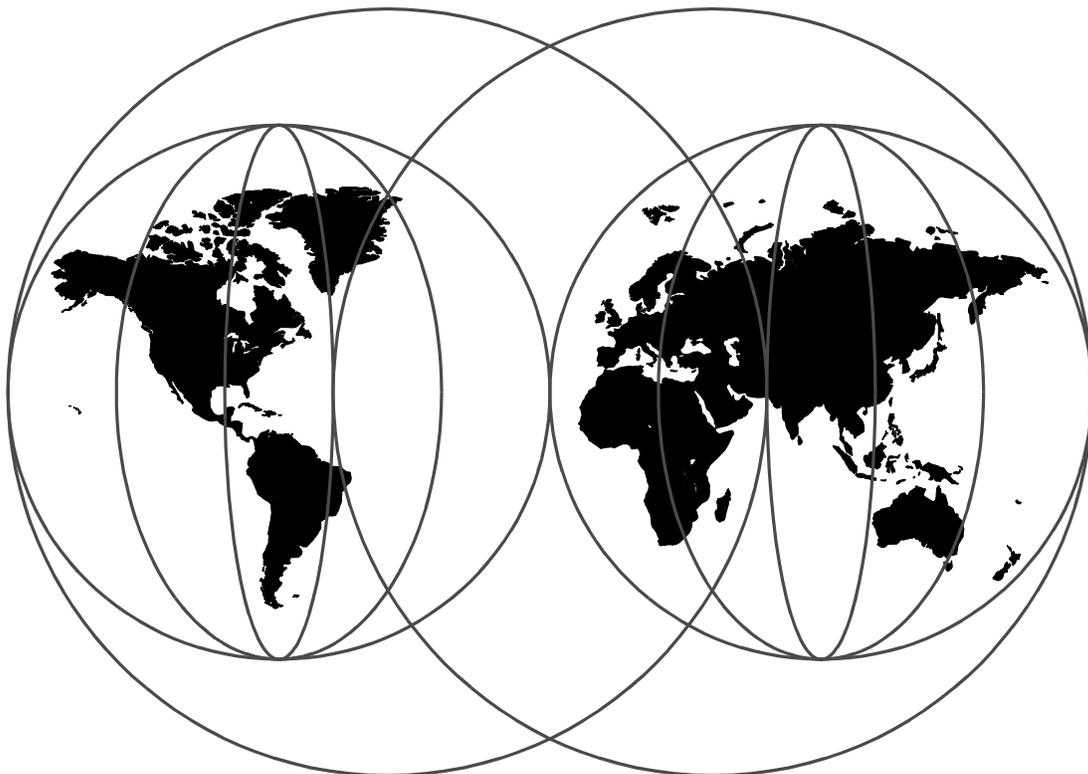


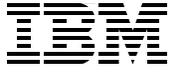
Getting Started with ADSM: A Practical Implementation Guide

*Arnold Balingit, Ross Battaglia, J.P Houle, Mathis Landzettel,
Armando Lemos da Silva Filho, Rod MacLeod, Andy Pattinson, Patrick Randall, Phil Thomas*



International Technical Support Organization

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



International Technical Support Organization

SG24-5416-00

Getting Started with ADSM: A Practical Implementation Guide

April 1999

Important: On-line Version Available!

This book-formatted version is valid only for the date of publication.

The most recent version of this material is maintained on the ITSO internet webpages at <http://www.redbooks.ibm.com/solutions/systemsmanagement/adsm/baseimpl.html>.

Take Note!

Before using this information and the product it supports, be sure to read the general information in Appendix C, "Special Notices" on page 257.

First Edition (April 1999)

This edition applies to Version 3 Release 1 of ADSTAR Distributed Storage Manager for AIX, program 5765-C43; Version 3 Release 1 of ADSTAR Distributed Storage Manager for HP-UX, program 5639-D92; Version 3 Release 1 of ADSTAR Distributed Storage Manager for MVS, program 5655-A30; Version 3 Release 1 of ADSTAR Distributed Storage Manager for Sun Solaris, program 5639-D91; Version 3 Release 1 of ADSTAR Distributed Storage Manager for Windows NT program 5639-C59.

Note:

IBM and Tivoli have announced the Tivoli ADSM V3 product, teaming along with the IBM ADSM V3. Their functionality is the same and for the purpose of this redbook, we refer to both products as ADSM. For information about IBM ADSM and Tivoli ADSM, please visit the IBM ADSM Web site at url <http://www.storage.ibm.com/software/adsm> and Tivoli ADSM Website at url <http://www.tivoli.com/adsm>.

Comments may be addressed to:
IBM Corporation, International Technical Support Organization
Dept. 471F Building 80-E2
650 Harry Road
San Jose, California 95120-6099

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

Contents

Figures	xi
Tables	xiii
Preface	xv
The Team That Wrote This Redbook	xv
Comments Welcome	xvii
Chapter 1. ADSM Implementation Checklists	1
1.1 The Big Picture	1
1.2 Redbook Support Material	2
1.3 Planning Checklist	5
1.4 Server Implementation Checklist	5
1.5 Client Implementation Checklist	7
1.6 Operations Checklist	8
Chapter 2. ADSM Implementation Planning	9
2.1 Client Environment Data	9
2.1.1 Client Name	10
2.1.2 Contact Information	10
2.1.3 Operating System	10
2.1.4 Total Storage Available	11
2.1.5 Total Storage Used	11
2.1.6 GB Changed per Backup	11
2.1.7 Number of Files Backed Up	11
2.1.8 Number of Backup Versions	11
2.1.9 Data Compression	12
2.1.10 Backup Window Times	13
2.1.11 Backup Number of Hours	13
2.1.12 Required Recovery Time Frame	13
2.1.13 ADSM Restore Time Frame	13
2.1.14 GB Copied per Archive	13
2.1.15 Number of Files Archived	13
2.1.16 Number of Archives Kept	14
2.1.17 Archive Frequency	14
2.1.18 Archive Window Times	14
2.1.19 Archive Number of Hours	14
2.1.20 Policy Domain	14
2.1.21 Client Option Set	14
2.2 Data Retention Requirements	15
2.2.1 Group Name	15
2.2.2 Number of Backup Versions	15
2.2.3 Backup File Retention Period	15
2.2.4 Number of Deleted Versions	15
2.2.5 Deleted File Retention Period	16
2.2.6 Offsite Copies	16
2.2.7 Onsite Collocation	16
2.2.8 Offsite Collocation	17
2.2.9 Archive Retention Period	17
2.3 Server Architecture Considerations	17
2.3.1 Server Platform	17

2.3.2	System Size	20
2.3.3	Multiple ADSM Servers	20
2.3.4	Server Licensed Features	21
2.3.5	Network	22
2.4	Disk Considerations and Sizing	23
2.4.1	The Disk Subsystem	23
2.4.2	ADSM Database	23
2.4.3	ADSM Recovery Log	26
2.4.4	Primary Disk Storage Pool	27
2.4.5	Device Configuration Table and Volume History File	28
2.4.6	Total Disk	28
2.5	Tape Considerations and Sizing	29
2.5.1	Tape Devices	30
2.5.2	To Library or Not to Library?	30
2.5.3	Number of Tape Drives	30
2.5.4	Number of Tape Volumes	31
2.5.5	MVS Tape Management	32
2.6	Administrator IDs	32
2.7	License Considerations	33
2.7.1	Server Type	34
2.7.2	Network Communications	34
2.7.3	Hierarchical Storage Manager	34
2.7.4	Backup-Archive Client	34
2.7.5	Advanced Device Support	34
2.7.6	Enterprise Administration	34
2.7.7	Disaster Recovery Manager	34
2.7.8	OpenSystems Environment	34
2.7.9	Server-to-Server Virtual Volumes	35
2.8	Other Considerations	35
Chapter 3. ADSM Server		37
3.1	Latest Code Updates	37
3.2	Code Installation	37
3.3	Customization	38
3.3.1	Options File	38
3.3.2	Settings	40
3.4	Server Implementation	42
3.4.1	Load Sample Scripts	42
3.4.2	Format Server Volumes	43
3.4.3	Label Tape Volumes	45
3.4.4	Windows NT Configuration Wizards	45
Chapter 4. Database and Recovery Log		47
4.1	Defining a New Database Volume	47
4.2	Defining a New Recovery Log Volume	49
4.3	Setting the Log Mode to ROLLFORWARD	51
4.4	Defining Database Backup Trigger	53
4.5	Setting the Expansion Trigger	53
4.5.1	Database Space Trigger	54
4.5.2	Recovery Log Space Trigger	55
4.5.3	When SPACETRIGGER Fails	55
4.6	Mirroring	56
4.6.1	Database Mirroring	57

4.6.2	Recovery Log Mirroring	58
4.7	Removing Default Volumes	60
4.7.1	Removing the Default Database Volume	60
4.7.2	Removing the Default Recovery Log Volume	61
4.8	Database Backup	62
4.9	Additional Commands	62
4.9.1	EXTEND DB	63
4.9.2	EXTEND LOG	63
4.9.3	REDUCE DB	64
4.9.4	REDUCE LOG	64
4.9.5	DELETE VOLHIST TYPE=DBBACKUP	64
Chapter 5.	Data Storage	67
5.1	Storage Pools	67
5.1.1	Primary Storage Pools	68
5.1.2	Copy Storage Pools	68
5.2	Storage Pool Hierarchy	68
5.3	Movement of Data between Storage Pools	69
5.3.1	Migration	70
5.3.2	MAXSIZE	70
5.4	Reclamation	71
5.4.1	Single Drive Reclamation	71
5.4.2	Reclamation of Offsite Volumes	72
5.5	Reduce Restore Times	72
5.5.1	Collocation	72
5.5.2	Disk Caching	73
5.6	RAID	74
5.7	Setting Up Storage Pools	75
5.7.1	Defining a Library — Physical	77
5.7.2	Defining a Drive — Physical	77
5.7.3	Defining a Device Class — Logical	78
5.7.4	Defining a Storage Pool — Logical	80
5.7.5	Defining Storage Pool Volumes - Logical.	83
5.8	Recommended Setup	85
5.8.1	Defining an Automated Library	85
5.8.2	Defining a Drive in an Automated Library	85
5.8.3	Defining a Device Class to an Automated Library	86
5.8.4	Defining Storage Pools	88
5.8.5	Defining Storage Pool Volumes	90
5.8.6	Deleting Default Storage Pools	93
5.9	Additional Commands	94
5.9.1	Audit Library	94
5.9.2	Audit Volume.	94
5.9.3	Back Up Storage Pool	95
5.9.4	Checkin Libvolume	96
5.9.5	Checkout Library Volumes	96
5.9.6	Delete a Library.	97
5.9.7	Dismount a Volume	99
5.9.8	Migration	100
5.9.9	Move Data.	101
5.9.10	Query Content.	101
5.9.11	Query Occupancy	102
5.9.12	Rename a Storage Pool	102

5.9.13 SQL Commands	103
Chapter 6. Data Storage Policy	105
6.1 Setting Up Data Storage Policies	106
6.1.1 Defining a Policy Domain	107
6.1.2 Defining a Policy Set	108
6.1.3 Defining Management Classes	108
6.1.4 Defining Backup Copy Groups	110
6.1.5 Defining an Archive Copy Group	113
6.2 Recommended Setup	114
6.2.1 Defining Policy Domains	115
6.2.2 Defining Policy Sets	115
6.2.3 Defining Management Classes	115
6.2.4 Defining Backup Copy Groups	116
6.2.5 Defining the Archive Copy Group	117
6.2.6 Deleting the STANDARD Policy Domain	118
6.3 Verifying Policy Definitions	119
6.3.1 Backup Copy Groups	119
6.3.2 Archive Copy Groups	120
6.4 Validating and Activating a Policy Set	120
6.4.1 Validating the Recommended Policy Sets	120
6.4.2 Activating the Recommended Policy Sets	121
6.5 Enforcing Your Policy	121
Chapter 7. User Management	123
7.1 Administrators	123
7.1.1 Considerations	123
7.1.2 Default Environment	123
7.1.3 Recommended Administrators	124
7.1.4 Working with Administrators	126
7.2 Client Nodes	127
7.2.1 Considerations	127
7.2.2 Default Environment	128
7.2.3 Working with Client Nodes	128
7.3 Client Option Sets	132
7.3.1 Considerations	132
7.3.2 Default Environment	133
7.3.3 Recommended Client Option Sets	133
7.3.4 Associating a Client Node with a Client Option Set	133
7.3.5 Working with Client Option Sets	134
Chapter 8. Licensing Your ADSM System	137
8.1 Licensed Features	137
8.2 Registering Licensed Features	138
8.3 Saving Your Licenses	139
8.4 License Compliance	139
8.5 Monitoring Licenses	139
8.5.1 Displaying License Information	139
8.5.2 Auditing Licenses	140
8.5.3 Scheduling Automatic License Audits	140
8.6 ADSM Evaluation License	141
Chapter 9. Administrative Client	143
9.1 Code Installation	144

9.1.1	Administrative Client Code	144
9.1.2	Web Administrative Interface	144
9.2	Customization	144
9.2.1	Administrative Client	144
9.2.2	Web Administrative Interface	145
9.3	Interfaces	145
9.3.1	Command Line Interface	145
9.3.2	Web Administrative Interface	146
Chapter 10.	Backup-Archive Client	151
10.1	Code Installation	151
10.1.1	Backup-Archive Client Code	151
10.1.2	Web Backup-Archive Client	151
10.2	Customization	152
10.2.1	Backup-Archive Client	152
10.2.2	Web Backup-Archive Client	156
10.3	Interfaces	156
10.3.1	Command Line	157
10.3.2	GUI	167
10.4	Client Scheduler	174
10.4.1	Starting the Client Scheduler	175
10.4.2	Stopping the Scheduler	178
10.5	Web Client Usage	179
10.5.1	Remote Client GUI	179
10.5.2	Client Functions	179
10.5.3	Access Authorization	179
10.5.4	Starting the Web Client	180
10.5.5	Stopping the Remote Agent Services	182
Chapter 11.	Scheduling	183
11.1	The Wheel of Life	183
11.2	Administrative Schedules	184
11.2.1	Defining an Offsite Backup Schedule	186
11.2.2	Defining the Volume History Schedule	188
11.2.3	Defining a Migration Schedule	189
11.2.4	Defining a Reclamation Schedule	190
11.2.5	Defining an Expiration Schedule	191
11.2.6	Defining a Licensing Audit Schedule	191
11.3	Client Schedules	192
11.3.1	Defining a Client Schedule	192
11.3.2	Associating a Client with a Schedule	193
11.3.3	Verifying the Client Schedules	194
Chapter 12.	Operations	195
12.1	Server Procedures	195
12.1.1	Starting the ADSM Server	195
12.1.2	Stopping the ADSM Server	199
12.2	Event Monitoring	202
12.3	Daily Sanity Checks	204
12.3.1	Database and Recovery Log	204
12.3.2	Data Storage	207
12.3.3	Client-Server Activity	210
12.4	Tape Management	213
12.4.1	Labelling Tapes	215

12.4.2	Onsite Tape Management to Offsite	217
12.4.3	Offsite Tape Management to Onsite	221
12.4.4	Reclaiming Offsite Tapes	223
12.4.5	Database Backup Management	223
12.5	Error Conditions	226
12.5.1	ADSM Errors	226
12.5.2	Machine Errors	226
Chapter 13.	Performance Considerations	231
13.1	ADSM Server.	231
13.1.1	BUFPOOLSIZE	231
13.1.2	EXPINTERVAL	232
13.1.3	LOGPOOLSIZE	233
13.1.4	MAXSESSION	233
13.1.5	MOVEBATCHSIZE and MOVESIZETHRESH.	233
13.1.6	USELARGEBUFFER	234
13.1.7	TCPWINDOWSIZE	234
13.1.8	TXNGROUPMAX	235
13.2	Client Node	235
13.2.1	COMPRESSION	235
13.2.2	COMPRESSALWAYS	236
13.2.3	DIRMC	236
13.2.4	QUIET	236
13.2.5	LARGECOMMBUFFERS	236
13.2.6	TCPBUFFSIZE	236
13.2.7	TCPWINDOWSIZE	236
13.2.8	TXNBYTELIMIT	237
13.3	How to Measure Performance	237
Appendix A.	Planning and Sizing Worksheets	239
Appendix B.	Redbook Support Material	243
B.1	Macros	243
B.1.1	Define Administrators.	243
B.1.2	Define Client Option Sets.	244
B.1.3	Define Policy Structure	245
B.1.4	Define Schedules.	246
B.1.5	Define Server Scripts.	247
B.1.6	Create Storage Pools.	248
B.1.7	Delete Default Storage Pools.	249
B.2	Server Options Files.	249
B.2.1	AIX	250
B.2.2	MVS	251
B.2.3	Windows NT.	252
B.3	Client Options Files	254
B.3.1	AIX	254
B.3.2	NetWare.	255
B.3.3	Windows 95/98/NT.	256
Appendix C.	Special Notices	257
Appendix D.	Related Publications	259
D.1	International Technical Support Organization Publications	259
ADSM Redbooks		259

D.2 Redbooks on CD-ROMs	260
D.3 ADSM Product Publications	260
D.4 ADSM Online Product Library	261
How to Get ITSO Redbooks	263
IBM Redbook Fax Order Form	264
Index	265
ITSO Redbook Evaluation	275

Figures

1. Redbook ADSM Environment	2
2. ADSM Database and Its Volumes	48
3. ADSM Recovery Log and Its Volumes	50
4. ADSM Mirror View	56
5. ADSM Storage Pool and Its Volumes	67
6. Hierarchical Arrangement of Different Storage Devices	69
7. Migration between Storage Pools	70
8. Space Reclamation of Tape Volumes	71
9. Storage Pool Collocation on a Client Level	73
10. Device Support Worksheet from ADSM Web Site	86
11. Storage Pool Hierarchy for Our Recommended Setup	88
12. Policy Relationships and Resources	105
13. Data Storage Policy Components	106
14. Include Option Example	109
15. ARCHMC Example	109
16. DIRMC Example	109
17. Data Flow through Copy Groups	111
18. Recommended Policy Definitions	114
19. Web Administrative Interface Logon Screen	147
20. ADSM Enterprise Console in Web Administrative Interface	149
21. Web Backup-Archive Client Main Window	151
22. Backup-Archive Client GUI: Windows	167
23. Backup-Archive Client GUI: UNIX	168
24. GUI Backup Window	168
25. GUI Restore Window	169
26. GUI Archive Window	172
27. GUI Retrieve Window	173
28. Batch File for Scheduler Start on Windows	176
29. Web Backup-Archive Client Main Window	181
30. Web Client Login Dialog Box	181
31. Scheduling of Operations	184
32. Administrative Commands to Back Up Storage Pools and Database	186
33. Script File for Offsite Backups	187
34. Launch Server Console in Foreground	198
35. Onsite and Offsite Distinction	213
36. ADSM Tape Processing	214
37. How ADSM Tapes Are Processed	215
38. ADSM DRM Tape Lifecycle	215
39. ADSM Error Entries in Windows Event Viewer	228
40. Detailed Event Information of ADSM Error	229

Tables

1. Redbook Support Material	3
2. Planning Checklist	5
3. Server Implementation Checklist	6
4. Client Implementation Checklist: Server Tasks	7
5. Client Implementation Checklist: Client Tasks	7
6. Daily Operations Checklist	8
7. Client Requirements Worksheet	9
8. Rules of Thumb for Selecting Percentage of Data Changed	11
9. Rules of Thumb for Selecting Data Compression	12
10. Typical Data Compression Ratios	12
11. Storage Policy Requirements Worksheet	15
12. Factors Affecting Collocation	16
13. ADSM Server Platform Considerations	18
14. Rough Estimates of Maximum Network Throughputs	22
15. Database Worksheet	26
16. Recovery Log Worksheet	26
17. Primary Storage Pool Worksheet	28
18. Device Configuration and Volume History Worksheets	28
19. Total ADSM Disk Required Worksheet	29
20. Tape Drive Configuration Worksheet	29
21. Administrator IDs Worksheet	33
22. License Requirements Worksheet	33
23. Server Options File Location	39
24. Server Options: Site Dependent	39
25. Server Options: Performance	40
26. Security Related Settings	41
27. Database and Recovery Log Volumes Allocation	57
28. Guide to Setting Up Storage Pools	76
29. Client Requirements Worksheet: Example	128
30. ADSM Licensed Features	137
31. Administrative Client Options	145
32. Web Client Options	156
33. Backup-Archive Client Interfaces by Platform	156
34. Backup Command Examples: UNIX	159
35. Backup Command Examples: Windows	159
36. Restore Command Examples: UNIX	161
37. Restore Command Examples: Windows	162
38. Archive Command Examples: UNIX	163
39. Archive Command Examples: Windows	163
40. Retrieve Command Examples: UNIX	165
41. Retrieve Command Examples: Windows	165
42. Recommended Schedules	185
43. Recommended Values for the BUFFPOOLSIZE	232
44. Client Requirements Worksheet	239
45. Storage Policy Requirements Worksheet	240
46. Database Worksheet	240
47. Recovery Log Worksheet	240
48. Primary Storage Pool Worksheet	241
49. Device Configuration and Volume History Worksheet	241
50. Total ADSM Disk Required Worksheet	241

51. Tape Drive Configuration Worksheet	242
52. Administrator IDs Worksheet	242
53. License Requirement Worksheet	242

Preface

This redbook provides material describing how to implement and operate ADSTAR Distributed Storage Manager (ADSM) Version 3 for people who have a general understanding of ADSM but have no practical experience of ADSM implementations.

Support materials in the form of example option files and macros are included where appropriate, and checklists lead you step-by-step through an implementation of the basic functions of ADSM:

- Planning and sizing
- Product implementation
- Customization
- Operation
- Performance and tuning

This practical guide is intended for the following audience:

- New Business Partners and other implementers aiming for ADSM Certification.
- System administrators, new to ADSM, who are asked to commence a basic ADSM implementation for the very first time
- ADSM administrators who want to learn more about the basic ADSM components and their implementation

The Team That Wrote This Redbook

This redbook was produced by a team of specialists from around the world working at the International Technical Support Organization San Jose Center.

Arnold Balingit is an IT Specialist in IBM Philippines. He has 7 years experience in the IT field. He holds a degree in Computer Engineering from Adamson University. His areas of expertise include ADSM implementation on a variety of supported platforms, project management, enterprise wide implementation of Systems Management Solutions (Tivoli), and development of Disaster Recovery Plans (DRPs).

Ross Battaglia is the Director of Technical Services at ETI, in Boca Raton, Florida, an IBM Business Partner, and Developer for the Tandem ADSM client. He has 20 years of experience in the data processing field. He holds a degree in Computer Science from Northeastern University. His areas of expertise include MVS systems programming, Storage Solutions, ADSM, Disaster Recovery, and Data Communications. He has written extensively on product evaluations and third party software.

J.P. Houle is a Senior IT Specialist with IBM Canada in Ottawa, Ontario. He is IBM Certified in AIX, ADSM and Tivoli Framework products and has been working in the UNIX field for 8 years. He holds a Bachelor's Degree in Mechanical Engineering with a concentration in Computer Integrated Manufacturing from Carleton University in Ottawa, Ontario, Canada.

Mathis Landzettel is a Project Leader at the International Technical Support Organization, San Jose Center. He joined IBM in 1994 after completing his degree in mathematics at the Technical University of Darmstadt. He writes extensively and teaches IBM classes worldwide on all areas of ADSM. Before joining the ITSO in 1998, Mathis worked in the ADSM development department in Mainz, Germany, as a software test team leader.

Armando Lemos da Silva Filho is the Information Security Officer at Deutsche Bank Brazil and also is an IBM certified specialist in ADSM. He has 11 years of experience in the data processing field. He has been working with ADSM, AIX, and NT for the past three years. His previous job experience includes Software Engineering and Technical Support. He holds a degree in Computer Science from PUC-SP. His areas of expertise include Hypertext Technology, Computer Security, and Operating Systems.

Rod MacLeod is an MVS Team Leader for TELUS Communications in Alberta, Canada. He has 19 years of experience in the data processing field. He holds a degree in Computation Science from the University of Saskatchewan. His areas of expertise include MVS systems programming, performance, capacity planning, and data storage. He has been working with ADSM for over three years. He has written articles and presented at conferences regarding his S/390 Parallel Sysplex implementation experiences.

Andy Pattinson is a Storage Consultant for Storm UK, an IBM business partner and European storage distributor in the north of England. He has 17 years experience in computer systems and storage. He is ADSM and NT certified. His areas of expertise include storage planning and hardware configuration, in addition to ADSM implementation.

Patrick Randall is a Distributed Storage Software Specialist at the International Technical Support Organization, San Jose Center. He has written eight redbooks on ADSM, teaches IBM classes worldwide on all areas of distributed storage, and is a consultant in Disaster and Business Recovery. Before joining the ITSO in July 1996, Patrick worked in IBM UK's Business Recovery Services as a Solutions Architect.

Phil Thomas is a senior Technical consultant for Aspect Computing, an IBM Business Partner in Australia. He has 22 years of experience in the Information Technology field. He holds a degree in Science from Sydney University and is an ADSM certified specialist. Apart from ADSM solution design and implementation, his areas of expertise include AIX, MVS, systems design, and Assembler programming. He has taught extensively in Asia and Australia.

Thanks to the following people for their invaluable contributions to this project:

Joan Bow
International Technical Support Organization, San Jose Center

Yvonne Lyon
International Technical Support Organization, San Jose Center

Mike Kaczmariski
IBM Storage Systems Division, Tucson

Comments Welcome

Your comments are important to us!

We want our redbooks to be as helpful as possible. Please send us your comments about this or other redbooks in one of the following ways:

- Fax the evaluation form found in “ITSO Redbook Evaluation” on page 275 to the fax number shown on the form.
- Use the online evaluation form found at <http://www.redbooks.ibm.com/>
- Send your comments in an internet note to redbook@us.ibm.com

Chapter 1. ADSM Implementation Checklists

The objective of this redbook is to provide material describing how to implement and operate ADSTAR Distributed Storage Manager (ADSM) Version 3 for people new to ADSM.

This chapter provides an overview of our redbook ADSM environment and implementation checklists for planning, installing, and operating that environment.

Our environment provides an integrated solution which incorporates complementary client and server options, basic performance recommendations, and operational processes. In our experience, this environment has been shown to satisfy the most common customer requirements and forms a sound basis for extension, or can be used as-is.

The checklists provide step-by-step processes to plan and implement an ADSM environment. Although geared towards our redbook environment, the checklists can be used for any implementation. There are separate checklists for planning the environment, server implementation, client implementation, and daily operations.

We have provided planning sheets, option files, and administrative macros to help plan and implement your ADSM environment. Appendix A, "Planning and Sizing Worksheets" on page 239 and Appendix B, "Redbook Support Material" on page 243 provide copies of those materials.

1.1 The Big Picture

Any ADSM solution consists of a number of pieces crafted to satisfy a particular set of requirements. These solution pieces include definitions for data storage management, policy management, user management, and operational management.

The difficulty is in determining how to craft each of these pieces to complete the solution jigsaw. This is complicated by the vast number of options and permutations that are possible with ADSM. We can help you.

We have developed a functional ADSM environment which has been shown to satisfy a number of key customer requirements. Those key requirements are:

- Multiple backup copies of files to be kept
- Second copy of backup data to be kept offsite
- Restore time to be minimized
- High level of automation

Our environment also incorporates basic performance tuning recommendations and operational procedures for onsite-offsite tape movement. It forms a sound platform for further expansion.

Figure 1 on page 2 shows the data storage management perspective of our ADSM environment. The figure shows the flow of data to and from the onsite storage pools and offsite copy pools. Some key features of this environment are:

- Separate storage pools for client directory information and client data
- Client data written to a disk storage pool, then migrated to tape storage pool
- Duplicate copies of onsite data created for offsite storage
- Mirrored ADSM database and recovery log

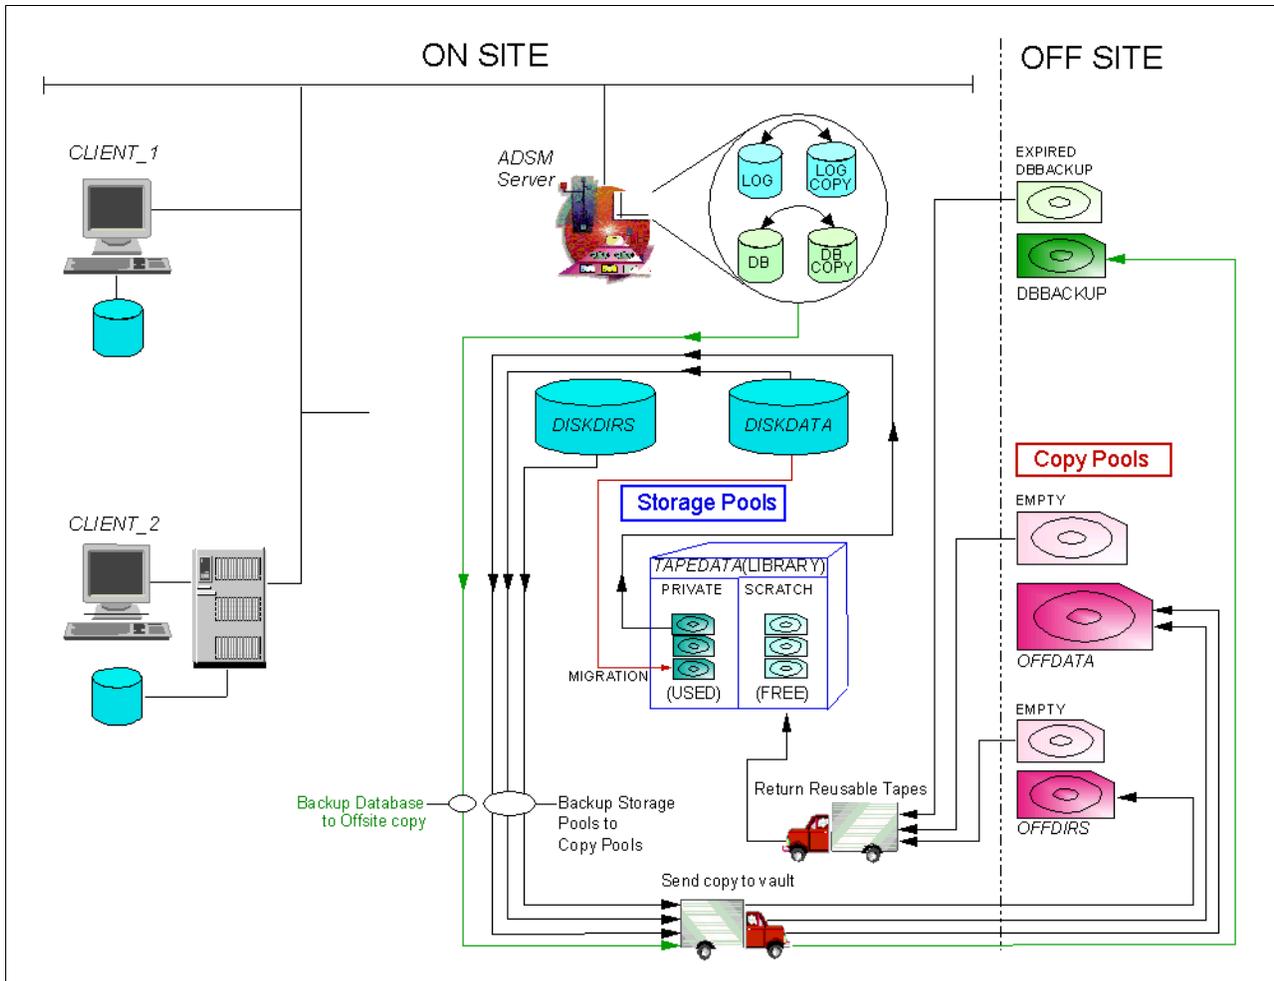


Figure 1. Redbook ADSM Environment

The two primary disk storage pools hold client directory information (DISKDIRS) and client data (DISKDATA). The remaining storage pool (TAPEDATA) is on tape and holds only client data. A copy of client data is stored in one offsite storage pool (OFFDATA), and a copy of the client directory data is stored in another storage pool (OFFDIRS).

1.2 Redbook Support Material

We provide worksheets, option files, and administrative macros to help plan and implement your ADSM environment.

The worksheets are used during the planning phase of an ADSM implementation which we use and discuss in Chapter 2, “ADSM Implementation Planning” on page 9.

The option files are examples to customize your ADSM server and clients. If you want to use them in your implementation, you need to modify the example files to fit into your environment, and then replace your options files with them.

The administrative macros are used to ease some steps of your ADSM implementation. Again, you need to modify some of the macros to meet specific requirements in your ADSM environment.

The redbook support material is available in softcopy on the internet from the redbooks Web server. Point your Web browser to:

<ftp://www.redbooks.ibm.com/redbooks/SG245416>

Alternatively you can go to:

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

and select Additional Redbook Materials (or follow the instructions given, since the Web pages change frequently!).

We provide two files. Each file contains all the support material in compressed format. The UNIX platform file is named `sg245416.tar`, and the Windows platform file is named `sg245416.zip`.

Table 1 on page 3 lists the contents of the support material from those files.

Table 1. Redbook Support Material

File name	Contents
<code>dsmopt.aix</code>	Client user options file for AIX
<code>dsmopt.nw</code>	Client options file for NetWare
<code>dsmopt.win</code>	Client options file for Windows 95/NT
<code>dsmserv.aix</code>	Server options file for AIX
<code>dsmserv.mvs</code>	Server options file for MVS
<code>dsmserv.win</code>	Server options file for Windows NT
<code>dsmsys.aix</code>	Client system options file for AIX
<code>mac.admins</code>	Administrative macro to define administrators
<code>mac.optionsets</code>	Administrative macro to define client option sets
<code>mac.schedules</code>	Administrative macro to define administrative and client schedules
<code>mac.scripts</code>	Administrative macro to define server scripts
<code>mac.stgcreate</code>	Administrative macro to create storage pools
<code>mac.stgdelete</code>	Administrative macro to delete default storage pools
<code>mac.policy</code>	Administrative macro to define policy domains, policy sets, management classes, and copy groups.

File name	Contents
planning.123	Planning spreadsheets (Lotus 123 format)
planning.xlw	Planning spreadsheets (Microsoft Excel format)
readme.1st	Contents of Support Materials

We recommend that you download the support material files into a separate directory on a system from which you can run an administrative command client. In our experience, the implementation works best when you choose a UNIX or Windows platform for that system.

If you want to use our redbook support material, you have to load the macro files into ADSM. We assume the administrative client is running in a Windows environment and the macro files are located in directory C:\redbook. You must invoke the administrative command line client with the ITEMCOMMIT parameter, otherwise the macro command will fail with a series of error messages.

The following command shows how to load the macro mac.admins to create the recommended administrative IDs. To load other macros, you need to specify another macro file.

```
C:\Program Files\IBM\ADSM\saclient>dsmdmc -itemcommit
ADSTAR Distributed Storage Manager
Command Line Administrative Interface - Version 3, Release 1, Level 0.6
(C) Copyright IBM Corporation, 1990, 1997, All Rights Reserved.

Enter your user id: admin

Enter your password: ****

Session established with server PAGOPAGO: Windows NT
Server Version 3, Release 1, Level 2.13
Server date/time: 02-02-1999 17:02:09 Last access: 02-02-1999 17:01:55

adsm> macro c:\redbook\mac.admins

ANS8000I Server command: 'register admin sysadmin sysadmin contact='System Administrator''
ANR2068I Administrator SYSADMIN registered.
ANS8000I Server command: 'grant authority sysadmin classes=system'
ANR2076I System privilege granted to administrator SYSADMIN.
ANS8000I Server command: 'register admin support support contact='System Support''
ANR2068I Administrator SUPPORT registered.
ANS8000I Server command: 'grant authority support classes=system'
ANR2076I System privilege granted to administrator SUPPORT.
ANS8000I Server command: 'register admin reporter reporter contact='System Reporting''
ANR2068I Administrator REPORTER registered.
ANS8000I Server command: 'register admin helpdesk helpdesk contact='Client Administrator''
ANR2068I Administrator HELPDESK registered.
ANS8000I Server command: 'query admin'
```

Administrator Name	Days Since Last Access	Days Since Password Set	Locked?	Privilege Classes
ADMIN	<1	17	No	System
HELPDESK	<1	<1	No	
REPORTER	<1	<1	No	
SERVER_CONSOLE			No	System
SUPPORT	<1	<1	No	System
SYSADMIN	<1	<1	No	System

1.3 Planning Checklist

Proper planning for your ADSM environment is critical to the success of the implementation. You should review all of Chapter 2, “ADSM Implementation Planning” on page 9.

The tasks contained in the ADSM planning checklist are shown in Table 2 on page 5. You should complete all these tasks before implementing the ADSM environment.

Table 2. Planning Checklist

✓	Tasks	Refer to
	Download redbook support materials	1.2 on page 2
	Complete client requirements worksheet	2.1 on page 9
	Complete data retention worksheet	2.2 on page 15
	Choose server platform	2.3.1 on page 17
	Size ADSM server	2.3.2 on page 20
	Complete licensing worksheet	2.3.4 on page 21
	Determine network load	2.3.5 on page 22
	Complete ADSM database worksheet	2.4.2 on page 23
	Complete ADSM recovery log worksheet	2.4.3 on page 26
	Complete ADSM storage pool worksheet	2.4.4 on page 27
	Complete ADSM disk worksheet	2.4.6 on page 28
	Determine tape library	2.5.2 on page 30
	Determine number of tape drives	2.5.3 on page 30
	Determine number of tape volumes	2.5.4 on page 31
	Complete administrator worksheet	2.5.5 on page 32

1.4 Server Implementation Checklist

The server checklist identifies those tasks you must complete to set up the redbook ADSM server environment. The tasks contained in the checklist are shown in Table 3 on page 6. These tasks are performed by either the system administrator or ADSM administrator.

The checklist consists of a series of tasks which must be performed sequentially. Each task in the table has a reference to another section in this redbook. The referred section contains the specific details on how to complete that task. For some tasks we additionally refer to the macro file we provide as described in 1.2, “Redbook Support Material” on page 2.

Table 3. Server Implementation Checklist

✓	Tasks	Refer to	Macro
	Download latest server code fixes	3.1 on page 37	
	Install server code	3.2 on page 37	
	Update server options file	3.3.1 on page 38	
	Install Web administrative client	9.1.2 on page 144	
	Create database volumes	4.1 on page 47	
	Create recovery log volumes	4.2 on page 49	
	Mirror database	4.6.1 on page 57	
	Mirror recovery log	4.6.2 on page 58	
	Remove default database volumes	4.7.1 on page 60	
	Remove default recovery log volumes	4.7.2 on page 61	
	Set up server licensing	8.2 on page 138	
	Define tape libraries	5.8.1 on page 85	
	Define tape drives	5.8.2 on page 85	
	Define device classes	5.8.3 on page 86	
	Change server run-time settings	3.3.2 on page 40	
	Set recovery log mode	4.3 on page 51	
	Set up database trigger	4.4 on page 53	
	Set up database expansion trigger	4.5.1 on page 54	
	Set up recovery log expansion trigger	4.5.2 on page 55	
	Define storage pools	5.8.4 on page 88	stgcreate
	Define storage pool volumes	5.8.5 on page 90	
	Remove default storage pools	5.8.6 on page 93	stgdelete
	Define policy domains	6.2.1 on page 115	policy
	Define policy sets	6.2.2 on page 115	policy
	Define management classes	6.2.3 on page 115	policy
	Define backup copy groups	6.2.4 on page 116	policy
	Define archive copy groups	6.2.5 on page 117	policy
	Activate new policy	6.4.2 on page 121	
	Remove default policy management definitions	6.2.6 on page 118	
	Define administrator IDs	7.1.3 on page 124	admins
	Define administrative schedules	11.2 on page 184	schedules
	Define client schedules	11.3.1 on page 192	schedules
	Create client option sets	7.3.3 on page 133	optionsets

1.5 Client Implementation Checklist

The client implementation checklist consists of two parts that identify those tasks you must complete to set up the redbook ADSM client environment. The tasks contained in the checklists are shown in Table 4 on page 7 and Table 5 on page 7.

Each checklist consists of a series of tasks which must be performed sequentially. Each task in the table has a reference to another section in this redbook. The referred section contains the specific details on how to complete that task.

The first checklist consists of tasks performed at the ADSM server. These tasks are performed by the ADSM administrator.

Table 4. Client Implementation Checklist: Server Tasks

✓	Tasks	Refer to
	Register client node	7.2.3.1 on page 128
	Associate client nodes with schedules	11.3.2 on page 193
	Associate client nodes with client option set	7.3.4 on page 133
	Grant authority for Web client access	7.2.3.3 on page 129
	Define event recording	12.2 on page 202

The second checklist consists of tasks performed at the ADSM client. These tasks are performed by the administrator of that client system.

Table 5. Client Implementation Checklist: Client Tasks

✓	Tasks	Refer to
	Download latest client code	3.1 on page 37
	Install client code	10.1 on page 151
	Update environment	10.2.1.1 on page 152
	Update client options files	10.2.1.2 on page 153
	Start backup-archive client	10.3.2.1 on page 167
	Implement scheduler	10.4.1.2 on page 177
	Implement Web client	10.5.4 on page 180

1.6 Operations Checklist

The operations checklist consists of those tasks you should complete on a daily basis. The tasks contained in the checklist are shown in Table 6 on page 8. Each task in the table has a reference to another section in this redbook. The referred section contains the specific details on how to complete that task. This checklist does not include the tasks we recommend scheduling on a daily basis. See Chapter 11, “Scheduling” on page 183 for more information on scheduled operations.

Table 6. Daily Operations Checklist

✓	Tasks	Refer to
	Check completion of client and administrative events	12.3.3.2 on page 211
	Bring free offsite volumes back to onsite	12.4.3 on page 221
	Send copy tapes to offsite location	12.4.2 on page 217
	Send ADSM database copy to offsite	12.4.5 on page 223
	Check client restartable restores	12.3.3.4 on page 212

Chapter 2. ADSM Implementation Planning

A successful implementation of ADSM benefits enormously from planning prior to attempting to set up the environment.

In this chapter, we present a number of planning worksheets that lead you through gathering the client requirements and the data retention requirements in an orderly way. We assume you are somewhat familiar with ADSM concepts and terms.

We provide planning sheets, option files, server scripts, and administrative macros to help plan and implement your ADSM environment. Appendix B, "Redbook Support Material" on page 243 contains information about how to download those support materials and what those materials provide. Blank worksheets are provided in Appendix A, "Planning and Sizing Worksheets" on page 239.

The redbook support material is available in softcopy on the internet from the redbooks Web server. Point your Web browser to:

<ftp://www.redbooks.ibm.com/redbooks/SG245416>

Alternatively you can go to:

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

and select Additional Redbook Materials (or follow the instructions given, since the Web pages change frequently!).

2.1 Client Environment Data

ADSM exists to provide services to clients, so it makes sense to begin by gathering data about the client environment. An ADSM client is the machine from which ADSM backs up or archives data. It could be various flavors of a workstation, or a file server, or possibly a database server. Even though you may know the machine as a server, to ADSM it is a client. The ADSM server refers to the machine where the ADSM server code runs. The ADSM server stores and manages all the data backed up from ADSM clients.

Complete Table 7 on page 9 using a column for data from each client being considered for this ADSM implementation. Analyzing client data helps you make decisions about the ADSM server environment. The information collected in the table is used for node definitions and for calculating database, recovery log, and storage pool sizes.

This table is presented in a portrait orientation in the book due to space considerations. If you have a large number of nodes or use a spreadsheet version, you may find the table more workable by transposing it to a landscape orientation.

Table 7. Client Requirements Worksheet

	Client 1	Client 2	Client 3
Client name	MyPC	File Server 1	File Server 2
Contact information	Rod at 7992	Server Group	UNIX Group

	Client 1	Client 2	Client 3
Operating system	Windows 95	Windows NT	AIX 4.3
Total storage available (GB)	3	10	33
Total storage used (GB)	0.75	8	25
GB changed per backup	0.01	0.8	1.5
Number of files backed up	5000	15000	2.5 GB
Number of versions kept	2	2	7
Data compression	0.5	0.5	0.66
Backup window times	16:30 - 08:00	20:00 - 07:00	23:30 - 04:30
Backup number of hours	15.5	11	5
Required recovery time frame	48 hours	24 hours	10 hours
ADSM restore time frame	24 hours	6 hours	8 hours
GB copied per archive			5
Number of files archived			17500
Number of archives kept			12
Archive frequency			Monthly
Archive window times			Month end
Archive number of hours			24 hours
Policy domain	Workstn	Workstn	Server
Client option set	Windows	Windows	AIX

2.1.1 Client Name

Enter the name ADSM will use for each client. Each name must be unique. We recommend using the machine name for the ADSM client name, so various groups such as Help Desk personnel or end users can easily correlate the client node to the ADSM name without having to look in a translation table.

By allowing the ADSM client name to default to the machine name on Windows clients, you are able to roll out a standard ADSM options file to numerous clients automatically without having to modify the options file for each client.

2.1.2 Contact Information

Enter information identifying the contact person or group responsible for this client.

2.1.3 Operating System

Identify the operating system and level the client is using. Any clients using operating systems not supported by ADSM will have to be handled separately.

2.1.4 Total Storage Available

Calculate the total amount of usable disk storage in GB available on the client. This number is the amount of disk storage seen by the client file system and does not contain the actual amount installed and used by a RAID or mirroring implementation.

2.1.5 Total Storage Used

Calculate the total amount of disk storage in GB currently in use or expected to be used on the client. If this number is unknown, use the *Total storage available* from above.

2.1.6 GB Changed per Backup

Calculate the amount of storage, data, or files changed between backup cycles. It indicates how long the client will be busy backing up, and how robust a network is required to complete all the backups each backup cycle. It is used to estimate disk and tape storage requirements on the ADSM server. Table 8 on page 11 presents some typical percentages of data changed for various sizes and types of data.

If your data change percentage is known, then enter that number in the worksheet. Otherwise, use Table 8 on page 11 to estimate a rate, keeping in mind that your numbers may vary. A high estimate is better than an estimate that is too low.

Table 8. Rules of Thumb for Selecting Percentage of Data Changed

Configuration	Percentage of Data Changed (%)
Large, busy file server	10
Smaller, less busy file server	5
Workstation	1
Database using utilities or ADSMConnect Agent	10-20
Database not using utilities or ADSMConnect Agent	100

2.1.7 Number of Files Backed Up

Calculate the total number of files to be backed up for each client. It is used to estimate disk and tape storage requirements on the ADSM server.

If the number of files is unknown, two values are possible. Enter either an estimate of the number of files in this field, or enter 5% of the *Total Storage Used* field, in GB.

2.1.8 Number of Backup Versions

Determine the number of changed copies that you want to keep of a file that exists on the client when the backup task runs. How many different versions of that file do you want to be able to restore? For example, if backup runs every night and a file changes every day, and you want to be able to restore from one week ago, then you would choose 7 as the number of backups to keep.

2.1.9 Data Compression

Estimate a data compression rate. To do that, you must decide whether to use data compression on the client.

ADSM allows a client to compress data prior to sending it to the ADSM server. To decide whether to use ADSM compression, take into account the speed of your network, the speed of the client CPU, and whether the ADSM server storage devices are capable of compression. Table 9 on page 12 provides some criteria for selecting data compression.

Table 9. Rules of Thumb for Selecting Data Compression

Configuration	Compression
Dial-up connection	Yes
Your network is approaching capacity	Yes
ATM network	No
CPU model of the client is less than Pentium II 100 Mhz (or equivalent)	No
CPU model of the client is greater than Pentium II 100 Mhz (or equivalent)	Yes
No hardware compression on tape devices	Yes
Small capacity or slow response storage devices	Yes

Compression takes time, so you need to decide whether the compression is helping or hurting the total elapsed time of your operation. In general, fast networks such as ATM networks do not benefit from compression because network bandwidth is usually not saturated. Conversely, older client CPU models may run slowly because the CPU cannot compress data fast enough to keep the network connection busy.

Many sites have various combinations of these configurations, such as slow networks, slow CPUs and small devices. In this case, make a judgement call about using ADSM compression. Performing tests may show you the most efficient mode for your environment.

Data compresses to varying degrees depending on the content of the files. Data composed of text, or many repeated characters like blanks, compresses well. Data consisting of random characters like executables does not compress well, and may actually grow.

If you decide to use compression, enter your data reduction rate. Use Table 10 on page 12 to estimate the data reduction if your actual ratio is unknown.

If you decide not to use compression, or are unsure whether to use compression, enter one (no compression) in the *Data Compression* field.

Table 10. Typical Data Compression Ratios

	Compression Ratio	Data Reduction
Database data	3:1 - 4:1	0.66 - 0.75
Print and file server data	2:1	0.5

	Compression Ratio	Data Reduction
Executables, compressed data, encrypted data	1:1	1

2.1.10 Backup Window Times

Enter the times of the day between which ADSM must start and complete its backup cycle. This window is influenced by when client usage subsides, by availability requirements, and by network capacity usage time frames.

2.1.11 Backup Number of Hours

Calculate the number of hours in the backup window.

2.1.12 Required Recovery Time Frame

Identify the time frame in hours, which you have agreed to with the customer, to recover a client completely. This is the time from when the data is lost to the time the data is usable again. It includes the time to fix or replace the machine and the time to restore the data.

Recovery time frames are vitally important, as this is the whole reason you are backing up.

This field documents your service level agreement with the customer.

2.1.13 ADSM Restore Time Frame

Calculate the time frame, in hours, allotted to ADSM to restore (possibly all) the data to the client. This number is used to size factors affecting the restore process such as network throughput and the number of tape drives required.

To calculate the *ADSM Restore Time Frame*, subtract from the *Required Restore Time Frame*, above, the maximum time required to prepare the client machine for a restore. In reality this number will vary depending upon the complexity of the restore. In the worst case (a disaster), the time to prepare the machine may include contacting support, fixing or replacing the machine, installing an operating system, installing the ADSM client code, and connecting to the network. In the best case (just deleted data), the time to prepare the machine may be the same as the *Required Restore Time Frame*.

A full ADSM restore time frame typically takes significantly more time than a full ADSM backup time frame. We have observed restore times of 110% or more over a full backup.

2.1.14 GB Copied per Archive

Calculate or estimate the amount of data to be archived during each archive session. Archives target specific data files. Typically, you do not archive whole systems.

2.1.15 Number of Files Archived

Calculate or estimate the number of files archived in each archive session.

If the number of files is unknown, two values are possible. Enter either an estimate of the number of files in this field or enter 5% of the *GB copied per archive* field, in GB.

2.1.16 Number of Archives Kept

Identify how many archives will be kept. For example, if an archive is performed monthly, you may want to keep 12 copies of that archive.

2.1.17 Archive Frequency

Determine how often you want to perform an archive function. Archives are typically run less frequently than incremental backups. Time frames such as daily, weekly, monthly or yearly are common.

2.1.18 Archive Window Times

Enter the times between which ADSM must start and complete its archive cycle for this client. This window is influenced by when client usage subsides, by availability requirements, by network capacity usage time frames, or by when the data becomes available for archive.

Archive time frames may be less time sensitive than incremental backups if they are going against copies of data or against data already processed. Often it is not possible to automatically schedule this archive function using ADSM schedules, due to the dependency on other events external to ADSM.

This number may be a specific time frame such as between 23:00 and 04:00, or a more general time frame such as after month-end processing finishes.

2.1.19 Archive Number of Hours

Identify the number of hours available to complete the archive function.

2.1.20 Policy Domain

Leave this field blank. During the planning phase, this is unknown. We use this field later in Chapter 1.4, "Server Implementation Checklist" on page 5 to contain the policy domain chosen for this client. For example, if you are implementing our Redbook configuration, the policy domains would be `SERVER` or `WORKSTN`.

For more information on policy domains, see Chapter 6, "Data Storage Policy" on page 105.

2.1.21 Client Option Set

Leave this field blank. During the planning phase, this is unknown. We use this field later in Chapter 1.5, "Client Implementation Checklist" on page 7 to contain the client option set chosen for this client. For example, if you are implementing our Redbook configuration, the client options set names would be `AIX` or `WINDOWS` or `NETWARE`.

For more information on client option sets, see Chapter 7.3, "Client Option Sets" on page 132.

2.2 Data Retention Requirements

In this section we identify the requirements for managing the data received from the clients. Categorize your data into a small number of groups with similar requirements. This table provides information to create copy groups under Management Classes in ADSM, to calculate storage pool sizes, and to calculate the number of tapes required to hold the data.

Complete a column in Table 11 on page 15 for each different group. We show two example groups.

Table 11. Storage Policy Requirements Worksheet

	Example 1	Example 2
Group name	Workstn	Server
Number of backup versions	2	7
Backup file retention period	90	NOLIMIT
Number of deleted versions	1	2
Deleted file retention period	60	180
Offsite copies	Yes	Yes
Onsite collocation	No	Yes
Offsite collocation	No	No
Archive retention period	None	None

2.2.1 Group Name

Choose a descriptive name for each different categorized group of data. In our example, the *Workstn* group is used for files from a workstation, while the *Server* group is used for all data from a large file server.

2.2.2 Number of Backup Versions

Determine the number of changed copies that you want to keep of a file that exists on the client when the backup task runs. How many different versions of that file do you want to be able to restore? For example, if backup runs every night and a file changes every day, and you want to be able to restore from one week ago, then you would choose 7 as the number of backups to keep.

Use the field *Number of Backup Versions* as a basis to group your data.

2.2.3 Backup File Retention Period

Determine the number of days you want to keep a backup version of a file (other than the current version). There are two options: keep the backup version for a set number of days; or specify NOLIMIT, which implies that you want ADSM to retain all backup versions, other than the most recent version, indefinitely.

2.2.4 Number of Deleted Versions

Determine how many versions of a file to keep after the file has been deleted from the original file system. This parameter comes into force during the first

backup cycle after the file has been deleted. For example, assume you are keeping seven versions of a file as specified above, and you have set this parameter to one. When the next backup cycle runs after the file has been deleted off the client, ADSM will flag the six oldest backup versions of the file for deletion and keep the most current backup version.

2.2.5 Deleted File Retention Period

Determine the number of days you want to keep the backup versions of a file after it has been deleted from the client. There are two options: keep the backup version for a set number of days; or NOLIMIT, which implies that you want to keep the backup version indefinitely.

For example, if you are keeping one version of a deleted file, and you set this parameter to 60, then 60 days after this file is noticed by ADSM as being deleted from the client file system, the one remaining backup version will be deleted from ADSM.

2.2.6 Offsite Copies

Determine if you wish to send a copy of the data offsite. Copying data to a removable device like tape allows the data to be taken offsite. An offsite copy along with other procedures provides recoverability in the event that the ADSM server becomes unusable, or that data on the ADSM server becomes corrupted.

Enter Yes to use offsite copies or No to not use offsite copies.

2.2.7 Onsite Collocation

Determine if you wish to use onsite collocation.

ADSM uses collocation to dedicate as few tapes as required to hold all of one client's files. Collocation reduces elapsed time for multiple file restores and full client restores at the expense of using more tapes and taking more ADSM management time for migration and for storage pool copies.

Collocating by client allows as many clients to be restored simultaneously as you have tape drives. If you have stringent restore requirements, collocation and numerous tape drives makes sense.

Table 12 on page 16 highlights some factors affecting whether to use collocation.

Table 12. Factors Affecting Collocation

	Onsite Collocation	Offsite Collocation
Short restore window	Yes	Yes
Less than 10 Clients, each 1 GB storage	No	No
More than 50 clients, each 1 GB storage	No	No
More than 50 clients, each 20 GB or more	Yes	Yes/No
Limited disk or tape resources	No	No
Workstations	No	No
Database, print, and file servers	Yes	Yes

ADSM implementations that are small (small number of clients managing a small amount of data) do not see much benefit from collocation, due to the small number of tapes required for any restore. Large ADSM implementations often collocate both onsite and offsite files due to the amount of data required for the restore of a client.

Tape drive capacity is a consideration for collocation. To use the minimum number of tapes, each tape must be used to its maximum capacity. Smaller capacity tapes will tend to fill completely even with small clients. In this case, collocation may be useful. With large capacity tape devices and collocation, a small client may not be able to fill a tape. With a large number of clients, significant numbers of tape volumes may be required.

See 5.5.1, “Collocation” on page 72 for more information on collocation.

2.2.8 Offsite Collocation

Determine if you wish to use offsite collocation. See 2.2.7, “Onsite Collocation” on page 16 for considerations.

2.2.9 Archive Retention Period

Determine how long you want to keep a file that is archived. Many sites set up a limited number of data groups with standard archive retention periods, such as seven days, 31 days, 180 days, 365 days, or seven years. Nonstandard requests for archive retention periods are slotted into the next larger retention period group. This reduces management complexity at the expense of keeping some data longer than actually required. If every nonstandard request is honored, the number of groups quickly becomes unmanageable.

2.3 Server Architecture Considerations

Having gathered information on the total client environment, you can now make decisions about the architecture of the ADSM server environment. This section deals with issues related to the ADSM server.

2.3.1 Server Platform

An ADSM Version 3 server runs on several platforms. How do you choose one platform over another? With only minor differences, an ADSM server provides the same functionality on every platform. The major differences between ADSM server platforms pertain to capacity, cost, installation, operation, supported devices, and installed user base. Each of these factors is explained below. Table 13 on page 18 summarizes these considerations in choosing an ADSM server platform.

An additional ADSM server option is the IBM 3466 Network Storage Manager. It is an ADSM-in-a-black-box approach that hides much of the management of ADSM. The sections “Client Environment Data” on page 9 and “Data Retention Requirements” on page 15 may be of use in sizing this device. See *A Practical*

Guide to Network Storage Manager (SG24-2242) for further information about Network Storage Manager.

Table 13. ADSM Server Platform Considerations

	Windows NT	UNIX	MVS
Installed user base	Medium	AIX - Large; HP-UX, Sun Solaris - Small	Small
Cost	Low	Medium	High
Capacity	Low - Medium	Medium - High	Medium - Very high
Platform installation	Simple	Simple	Complex
Operation	Simple	Complex	Very complex
Supported devices	Many	Many	Fewer

2.3.1.1 Installed User Base

The number of ADSM servers installed for a particular platform is a consideration. IBM ships new functions to the popular ADSM server platforms (AIX, MVS, and Windows NT) first. Currently ADSM server platform updates on AS/400 are done once per year due to the AS/400 software release cycle.

The more popular a platform, the more the software is exercised leading to a lower probability that you will find a unique problem. More popular platforms have a larger base of knowledgeable people to call upon for advice and to draw from for new hires.

2.3.1.2 Cost

Cost is a very dynamic area to discuss in a static manual, so we shall only generalize! Check for special promotions and other discounts before committing to acquiring a particular platform configuration.

Cost is further subdivided into platform costs and ADSM software license costs. Platform costs include the cost to acquire all the hardware and software to run the platform exclusive of the ADSM software license. It ranges from very low for Windows NT to extremely high for MVS. The high cost of installing a new MVS platform usually precludes it from being selected to run an ADSM server exclusively.

ADSM license costs vary considerably, with Windows NT costs being the lowest, followed by AIX, HP-UX, Sun Solaris, and AS/400 costs. All of these are one-time charges to purchase the product. MVS licenses are available for a one-time charge or as a monthly license fee. You may wish to calculate the break-even point.

Be careful, when paying a one-time charge license fee, that a newer version of the software is not about to be released. If you wish to upgrade to this newer version quickly, you may be forced to pay significant one-time charges all over again. Licensing ADSM monthly on MVS does not suffer from this disadvantage.

ADSM Server costs include a license for only one client. To manage more clients, more client licenses must be purchased and registered on the server. The actual client code is free.

2.3.1.3 Capacity

The ADSM server has the capacity to manage a basically unlimited number of clients and an unlimited amount of data. The platform the ADSM server software is running on does limit what ADSM can administer. Various platforms have different capacities in regard to the CPU power it can deliver to ADSM, the number of devices it can attach, and the throughput it can deliver.

Choose your platform with growth in mind. Moving from a small platform to a larger platform of the same server type, such as from a small AIX box to a larger AIX box, is relatively simple. Starting at the top end of a server type and moving to another server type, such as from Window NT to AIX, involves exporting and importing each client separately. Although the procedure is straightforward, it can be time-consuming and procedure-intensive.

2.3.1.4 Platform Installation

Installation consists of the platform installation and the ADSM server code installation. Platform installation consists of hardware installation and configuration, and of operating system installation and configuration. Installation of each platform requires specialized knowledge that will not be covered here.

The ADSM server code installation varies by platform in the specifics, but generally follows a similar procedure. Installation on Windows NT can be easier due to the Windows NT Wizards that have been provided. Installation of the ADSM server on other platforms is not difficult for an administrator familiar with the platform.

2.3.1.5 Operation

Operation of a platform varies from almost completely automatic on Windows NT, to very complex on MVS, with the UNIX platforms in the middle. The availability of tools to assist in managing the operation of the various platforms is nearly opposite, in that the MVS environment has a rich, powerful assortment of tools, while Windows NT is lacking in this regard.

Operation of ADSM itself varies only in the way some operating system-specific ADSM commands are issued on each platform.

2.3.1.6 Supported Devices

There are a wide variety of supported devices on the Windows NT and UNIX platforms, including disk drives, tape drives, optical drives and automated tape libraries. MVS and AS/400 are limited to their standard choice of devices, but these devices generally have tremendous capacity. For a list of supported devices on each ADSM server platform, see the URL:

<http://www.storage.ibm.com/software/adsm/addevice.htm>

A concern with the smaller platforms is the ability to attach the required amount of devices as the environment grows. On larger platforms this concern is usually reduced.

2.3.1.7 Recommendations

If you already have an MVS or AS/400 system with capacity available, these platforms are a good choice. Otherwise, they are probably out of consideration due to their complexity, cost, or specialized nature.

If your site already runs HP-UX or Sun Solaris, these are acceptable choices.

If you have experience with Windows NT, and the ADSM implementation is small and will remain small, choose Windows NT.

For sites with little or no commitment to other platforms, choose AIX. It is robust enough to handle most implementations while remaining relatively cost effective. AIX-based ADSM systems are very well supported by IBM in both the operating system and in ADSM. There are large numbers of ADSM installations to gather knowledge from. The AIX platform (and others) scales from small to very large ADSM systems.

2.3.2 System Size

Choosing the correct platform CPU size and memory requirements is an inexact art. There is very little information to guide you in selecting an ADSM platform. As you would expect, the risk of choosing an inappropriate platform size increases as the size of the ADSM implementation increases. Small ADSM implementations are at less risk of choosing an incorrect platform size, and the incremental cost to scale up or down is small. Many sites start small and grow into larger systems. However, this is of little help if you are starting large.

ADSM is CPU-, I/O-, and memory-intensive. CPU is a function of the number of files to manage and how your platform processes I/O. A large number of small files is more CPU-intensive than a small number of large files. As the number of files and the amount of data to be moved increases, each backup, migration, storage pool copy, and expiration process will use more CPU to maintain the database entries. ADSM takes advantage of multiple processors. In our experience, MVS platforms seem to be more CPU-intensive than UNIX platforms. MVS sites should be aware that ADSM can use significant amounts of CPU. We have seen ADSM among the top five users of CPU on MVS systems.

I/O is the major part of ADSM processing. ADSM does very little else. Backups and restores are I/O intensive. Database updates and retrievals are I/O intensive. Storage pool management is I/O intensive. The I/O subsystem needs to be robust enough to handle this load. As the number of files and the amount of data climb, the need for a larger, faster I/O subsystem increases. Separate controllers or adapters for disk and tape devices become essential as the load increases.

Memory is used to cache database entries, among other things. As the number of files being managed increases (and thus the database size increases), the amount of memory that ADSM requires increases. MVS users note that ADSM likes to keep a large part of its address space in real storage. Since memory is relatively cheap, and you will never regret having too much memory, we recommend starting with 256 MB on most platforms.

2.3.3 Multiple ADSM Servers

When first setting up an ADSM environment, we recommend implementing a single ADSM server. Once experience has been gained, and the implementation

has grown enough to be reaching the capacity of the current server model family, a second server may be considered. Note that we recommend upgrading your current server hardware to its next larger model before considering a second server. ADSM can handle very large amounts of data or clients in one implementation. Currently, we have seen ADSM implementations with ADSM databases of 60 GB and larger (admittedly on large platforms).

2.3.3.1 Reasons to Consider Multiple Servers

Multiple ADSM servers can be configured to provide some redundancy and disaster recoverability in the event of an ADSM server outage. For example, a company with two sites A and B, may decide to install an ADSM system in each site. The ADSM system in site A would back up the data from site B and the ADSM system in site B would back up the data from site A. The loss of site A would mean that the ADSM system in site B (which holds the data for site A) could immediately start restoring data to spare equipment at site B. The site A ADSM (which was lost) could be recovered to the ADSM system at site B.

The Server-to-Server Communications feature, the Enterprise Management enhancements and the Enterprise Administration features make managing multiple servers easier by centralizing some administration functions and allowing changes to be replicated on some or all systems.

For very large or critical clients such as a large, enterprise-wide business intelligence complex, an ADSM server dedicated to the client may make sense.

In installations where network connectivity is slow or expensive, placing an ADSM server close to the client(s) may make sense. For example, for a business that has multiple file servers in each of a number of cities interconnected by a slow network, it may be appropriate to install an ADSM server in each city.

2.3.3.2 Disadvantages of Multiple Servers

Multiple servers increase costs. Two small server CPUs may be more expensive than one larger CPU of the same power. Where one automated tape library may be enough, multiple servers may require multiple automated libraries. Every ADSM server requires an ADSM server license. Some enterprise administration options also have to be licensed on each ADSM server to be managed.

Management of a multiple server environment is more complex, costly and time consuming than a single environment. Installation and maintenance procedures have to be repeated on each server. Confusion about where data is stored, in the event of a restore, may result. Some of these disadvantages can be reduced by using the Enterprise Administration feature.

2.3.4 Server Licensed Features

The base ADSM server license provides the backup and restore function and an archive and retrieve function for one client. Additional licenses are available for several clients, network connection, and advanced device support. Optional licenses can be acquired for the Hierarchical Storage Management (HSM) function, the Disaster Recovery Management (DRM) function, the Server-to-Server Virtual Volumes feature, and the Enterprise Administration feature. None of the last features are covered in this book, and we recommend that you become familiar with the basic ADSM operation prior to implementing these options.

2.3.5 Network

The network connection plays an integral part in providing service to the customer. It may consist of a combination of network interface cards, hubs, routers, gateways, wire, and software. ADSM server software, ADSM client software, the server platform and the client platform all have at least minimal monitoring or management capabilities. Often small networks have none of this, or the function may not be implemented or accessible. This makes the network a weak link in the overall management of the implementation.

Network design, implementation, and operation is beyond the scope of this book. However, we cover some basic recommendations.

2.3.5.1 Network Topology

Network topologies such as Ethernet, Token ring, FDDI, and ATM all work well with ADSM. Each has its strengths and weaknesses. In general, choose the fastest network topology you can afford.

To estimate the speed of the network required, use the following calculation.

Using Table 7 on page 9, calculate the total amount of data to be transferred during a backup window. For each client:

1. Multiply the *GB changed per backup* field by (1 - *Data compression* field).
2. Sum this number for all clients, giving the total data to be transferred.
3. Divide the total data transferred by the number of hours in your backup window to give the transfer rate in GB/hr.

To calculate archive speed, use the same calculation as above, but substitute *GB copied per archive* for *GB changed per backup*. If backups and archives will occur during the same time period, these two figures must be added together. Keep in mind that ADSM traffic likely will not be the only traffic on this network.

Use this number to decide what type of network speed is required to ensure that your backups finish in the backup window. Table 14 on page 22 provides some very rough estimates of network throughput for various network topologies, but bear in mind that as a network approaches capacity, throughput tends to drop substantially. The table provides only an indication of topologies that may or may not be capable of providing the required capacity. It is not meant to indicate actual throughput rates.

Table 14. Rough Estimates of Maximum Network Throughputs

Network Topology	MB/sec	GB/hr
10 Mb/sec Ethernet	1	3.6
100 Mb/sec Ethernet	10	36
4 Mb/sec Token Ring	0.4	1.4
16 Mb/sec Token Ring	1.6	5.8
FDDI	10	36
ATM 155 Mb/sec	15.5	55.8

Using the same methodology, work out a rough estimate of how long a restore will take. Use the *Total Storage Used* field instead of the *GB Changed per Backup*

field. This estimate is rough because ADSM management time and ADSM compression time (if used) are not considered.

2.3.5.2 Communication Protocol

Most network protocols such as TCP/IP and NETBIOS are supported by ADSM. TCP/IP is the most common communication method and possibly the easiest to set up from an ADSM perspective. Certain functions, such as server prompted mode and the Web clients, require TCP/IP.

2.3.5.3 Network Names

The ADSM server machine requires a name that clients can use to point to the ADSM server. If TCP/IP is used, create a Domain Name Server (DNS) entry for ADSM itself, as well as a DNS entry for the machine. For example, create DNS name ADSM1 for ADSM and ADSMSRV1 for the machine. If ADSM needs to be moved to another machine, only the DNS entry needs to be changed, instead of editing the ADSM options file for each client. Additionally, this allows for the addition of other ADSM servers. However, using DNS may impact backup availability if for some reason the service is down.

2.4 Disk Considerations and Sizing

ADSM requires disk to operate: it needs disk to hold the database, logs, and usually the primary storage pools. In this section we talk about how to choose a disk subsystem, and how to determine the correct amount of disk storage capacity to acquire.

2.4.1 The Disk Subsystem

In general, choose the fastest disk subsystem that you can afford. Slow disks may not be a hindrance for a small implementation, but as the site grows, disk access becomes a large percentage of the overall elapsed time to complete a task. Size the disk subsystem for growth because the vast majority of ADSM implementations grow substantially. Choose a disk model that meets your present estimated needs and has room for expansion.

See 5.6, "RAID" on page 74 chapter for a discussion about choosing RAID devices.

2.4.2 ADSM Database

The ADSM database size is based upon how many files are managed with ADSM, and whether the files are in a primary storage pool or a copy storage pool. The database holds two types of data: entries for backups and entries for archives. The database also holds items such as server scripts and the volume history, but typically, these are insignificant in sizing the database.

Backup Sizing calculates the size of the ADSM database holding backup entries. Archive Sizing calculates the size of the ADSM database holding archive entries. Use either or both depending on the data you will be storing in ADSM. If both are used, then add the calculated database sizes together to arrive at the total database size.

2.4.2.1 Backup Sizing

To estimate the ADSM database size for backup data, use the data collected in Table 7 on page 9 and Table 11 on page 15.

The calculation is based upon two types of numbers: the number of files backed up; and when the number of files backed up is not known, by a percentage of the data backed up. Steps 1, 2, 3, and 6 are based on the number of files. Steps 4 and 7 are based on the amount of data being backed up.

1. Sum the field *Number of Files Backed Up* for all clients, leaving out any fields containing estimates in GB.
2. Multiply this number by the *Number of Versions Kept*, giving a total number of files backed up.
3. Multiply this number by 600 bytes, to give bytes used for all known files backed up.
4. Sum of all the estimated files in GB. Take 5% of this number to give the estimated bytes for backed up files.
5. Add bytes for known files backed up to estimated bytes backed up to give total bytes for backed up files.
6. If copy storage pools are used (and they most likely will be used), multiply the total number of files backed up and calculated in Step 2 above by 200 bytes, giving the bytes for known copy storage pool files.
7. Sum all estimated files in GB. Take 1% of this number to give the estimated bytes for copy storage pool files.
8. Add bytes for known copy storage pool files to estimated bytes for copy storage pool files to give total bytes for copy storage pool files.
9. If offsite copies are used, multiply total bytes for copy storage pool files by 2, again giving total bytes for copy storage pool files.
10. Add the total bytes for backed up files to the total bytes for copy storage pool files to give the total bytes calculated for the database.
11. Calculate 135% of total bytes calculated for the database to give the database size. This is for overhead and for growth.

For example, using the sample data in Table 7 on page 9 and Table 11 on page 15, calculate the sample database size as follows:

1. $5,000 + 15,000 = 20,000$ files
2. $20,000 * 7 = 140,000$ files
3. $140,000 * 600 = 84,000,000$ bytes
4. $2,500 \text{ MB} * 0.05 = 125 \text{ MB}$
5. $84 \text{ MB} + 125 \text{ MB} = 209 \text{ MB}$ for backed up files
6. $140,000 * 200 = 28,000,000$ bytes
7. $2,500 \text{ MB} * 0.01 = 25 \text{ MB}$
8. $28 \text{ MB} + 25 \text{ MB} = 53 \text{ MB}$ for copy storage pool files
9. $53 \text{ MB} * 2 = 106 \text{ MB}$ (offsite is used)
10. $209 \text{ MB} + 106 \text{ MB} = 315 \text{ MB}$
11. $315 \text{ MB} * 1.35 = 350 \text{ MB}$ database size

2.4.2.2 Archive Sizing

To estimate the ADSM database size for archive data, use the data collected in Table 7 on page 9 and Table 11 on page 15.

The calculation is based upon two types of numbers: the number of files archived; and when the number of files archived is not known, by a percentage of the data archived. Steps 1, 2, 3, and 6 are based on the number of files. Steps 4 and 7 are based on the amount of data being backed up.

1. Multiply the *Number of files archived* by the *Number of archives kept*, giving a total number of files archived. Leave out any fields with estimate in GB.
2. Sum this number for all clients.
3. Multiply this number by 600 bytes to give bytes for all known files archived.
4. Sum all of the estimated files in GB. Take 5% of this number to give the estimated bytes for archived files.
5. Add bytes for known files archived to estimated bytes archived to give total bytes for archived files.
6. If copy storage pools are used (and they most likely will be used),
 1. Multiply the total number of files archived and calculated in Step 2 above by 200 bytes, giving the bytes for known copy storage pool files.
 2. Sum all estimated files in GB. Take 1% of this number to give the estimated bytes for copy storage pool files.
7. Add bytes for known copy storage pool files to estimated bytes for copy storage pool files to give total bytes for copy storage pool files.
8. If offsite copies are used, multiply total bytes for copy storage pool files by 2, again giving total bytes for copy storage pool files.
9. Add the total bytes for archived files to the total bytes for copy storage pool files to give the total bytes calculated for the database.
10. Calculate 135% of total bytes calculated for the database to give the database size for archives. This is for overhead and for growth.

2.4.2.3 Identify Databases

We recommend using the ADSM mirroring function for the database instead of a hardware mirror or operating system mirror, because ADSM has additional function to handle error conditions that may affect the mirrored copy.

If you are using ADSM mirroring, you need to plan for the mirror copy by doubling the amount of disk for the database.

Various file systems have different maximum capacities, so the database may have to be split across numerous volumes to make up your database size. We recommend that the primary and copy files for the database each be placed on a separate physical drive and controller if possible.

Complete Table 15 on page 26 with the database file names and the volume names for your primary database and copy database.

Table 15. Database Worksheet

Filename (Primary)	Volume	Size (MB)	Filename (Copy)	Volume	Size (MB)
/adsm/database/primary/file01	Vol1	175	/adsm/database/copy/file01	Vol1	175
/adsm/database/primary/file02	Vol2	175	/adsm/database/copy/file02	Vol2	175
Total		350	Total		350

2.4.3 ADSM Recovery Log

The size of the recovery log depends on the amount of data changed between ADSM database backups. The larger the amount of data, the larger the ADSM recovery log needs to be. Either a full ADSM database backup or an incremental ADSM database backup (in roll-forward mode) resets the ADSM recovery log back to empty. If the recovery log fills up completely, ADSM stops and you have to manually increase the size of the recovery log. This is painful and to be avoided at all costs.

To estimate the size of the recovery log, multiply the database size by the percentage of data that changes each backup cycle. Double this number to allow for two backup cycles to occur without a database backup. This gives a starting point for the recovery log.

For example, if the database size is 350 MB, and 5% of the data changes every backup cycle, then the estimated size for the recovery log would be $350 \text{ MB} \times 0.05 \times 2 = 35 \text{ MB}$.

We recommend using the ADSM mirroring function for the recovery log instead of a hardware mirror or operating system mirror, because ADSM has additional function to handle error conditions that may affect the mirrored copy.

If you are using ADSM mirroring, you need to plan for the mirror copy by doubling the amount of disk for the recovery log.

Various file systems have different maximum capacities, so the recovery log may have to be split across numerous volumes to make up your total recovery log size. We recommend that the primary and copy files for the recovery log each be placed on a separate physical drive and controller if possible.

Complete Table 16 on page 26 with the recovery log file names for your primary and copy recovery log.

Table 16. Recovery Log Worksheet

Filename (Primary)	Volume	Size (MB)	Filename (Copy)	Volume	Size (MB)
/adsm/log/primary/file01	Vol1	35	/adsm/log/copy/file01	Vol1	35
Total		35	Total		35

2.4.4 Primary Disk Storage Pool

When a client node backs up, archives, or migrates data, the data is stored in a primary storage pool.

To size the disk storage pool, calculate the amount of data that will be backed up during one backup cycle, transferred to the server, and stored on the server. This amount (plus a contingency for growth) is the recommended size the storage pool needs to be.

Primary storage pools (usually disk devices) can be made larger or smaller than the recommended size. A larger pool size would allow for more than one backup cycle of data to remain on disk, thus improving recall elapsed times. It allows for spare capacity for an unexpectedly large amount of backup data to prevent migration from running during backup. A smaller primary storage pool uses less disk, but then migration thresholds need be set and monitored at levels that keep the primary storage pools from filling up during the backup cycle (in order to avoid data spilling to the next storage pool).

We recommend using a primary disk storage pool at the recommended size to reduce interference from migration while backup is running.

To estimate a primary storage pool size, perform the following steps:

1. Using Table 7 on page 9, multiply the *GB changed per backup* by $(1 - \text{the Data compression rate})$ to give the total bytes transferred for each client.
2. Sum the total bytes transferred for all clients to give the total bytes transferred per backup cycle.
3. Add 15% to total bytes transferred per backup cycle to give the storage pool size. This allows for variability in the size and number of files per backup.

For example, using the sample figures in Table 7 on page 9, the *GB changed per backup* are 0.01, 0.8, and 1.5 while the *Data compression* figures are 0.5, 0.5, and 0.66 respectively.

1. Multiplying 0.01 by $(1 - 0.5)$ gives 0.005, and multiplying 0.8 by $(1 - 0.5)$ gives 0.4, and multiplying 1.5 by $(1 - 0.66)$ gives 0.5.
2. Summing 0.005, 0.4, and 0.5 gives 0.905 GB.
3. Adding 15% gives a storage pool size of 1.04 GB.

We recommend that the disk storage pools be allocated on fault tolerant hardware devices such as RAID 5 devices. If you are using mirroring, you need to plan for the mirror copy by doubling the amount of disk for the primary storage pool.

Various file systems have different maximum capacities, so the primary storage pool may have to be split across numerous volumes to make up your total primary storage pool size. We recommend that the disk storage pools be placed on their own disk devices and controller separate from the database and the recovery log, if possible.

Complete Table 17 on page 28 with the primary storage pool file names and volume names for your primary storage pool.

Table 17. Primary Storage Pool Worksheet

Filename	Volume	Size (MB)
/adsm/stgpool/file01	Vol1	200
/adsm/stgpool/file02	Vol2	200
/adsm/stgpool/file03	Vol3	200
/adsm/stgpool/file04	Vol4	200
/adsm/stgpool/file05	Vol5	200
/adsm/stgpool/file06	Vol6	200
Total		1,200

2.4.5 Device Configuration Table and Volume History File

The device configuration table and the volume history table also require disk, but typically, they are very small. The device configuration table contains entries for defined device classes and definitions for drives and libraries. Every volume that is used by ADSM is tracked in the volume history database, including the volume identifier for the database backups. The volume history information is periodically copied out to a volume history file that you can specify with the `VOLUMEHISTORY` option in the `dsmserv.opt` file.

We recommend that you have two copies of the device configuration table and the volume history file, in case one becomes unusable due to a software failure.

Complete Table 18 on page 28 with the device configuration and volume history filenames and sizes.

Table 18. Device Configuration and Volume History Worksheets

Name	Size (MB)
/adsm/devconfig1.out	0.5
/adsm/devconfig2.out	0.5
/adsm/volhist1.out	0.5
/adsm/volhist2.out	0.5
Total	2

2.4.6 Total Disk

Total disk refers only to the numbers discussed here. If you are using mirroring or some other version of RAID, you need to take that into consideration separately. The disk required to run the server platform operating system efficiently also has not been considered.

ADSM code requirements for disk vary depending on the server platform and release level. We use an estimate of 150 MB. Table 19 on page 29 summarizes the disk requirements for the ADSM server.

Table 19. Total ADSM Disk Required Worksheet

	Size (MB)
ADSM code (dependent on platform)	150
ADSM database	700
ADSM recovery log	70
Primary storage pools	1,200
Device configuration table and volume history table	2
Other (RAID, Operating system)	0
Total	2,122

2.5 Tape Considerations and Sizing

Most ADSM systems use tape as the medium for storing their data long term, such as in the copy pools. There are a variety of tape subsystems available for most ADSM server platforms. In this section, we discuss various tape configurations and how to size these configurations.

In general, choose the biggest, fastest, most automated tape drive solution you can afford.

Optical disks are an alternative to tape, but we recommend using tape for most ADSM implementations, as tape is faster and more convenient to work with.

Planning for and development of manual procedures in ADSM may be required to identify and move volumes from an offsite copy storage pool to the offsite storage location and back.

Complete the worksheet in Table 20 on page 29 using the information from the sections below. The data collected will be used when sizing tape libraries and then when defining them to ADSM.

Table 20. Tape Drive Configuration Worksheet

	Option
Library model	ASIC VLT DLT
Number of drives	3
Drive model	DLT
Number of onsite tape volumes	32
Number of offsite tape volumes	32
Number of database ba volumes	6
Number of scratch tapes	11
Total tape volumes required	81

2.5.1 Tape Devices

Tape drives come in 4mm, 8mm, DAT, 3570, 3590, and other device types. Each type of drive has a different data capacity, performance, cost, and reliability characteristics. Although data capacity and cost per megabyte stored are important, reliability is much more important. Having saved money buying tapes is small consolation when you are unable to restore your customer billing database due to a tape error.

In general, tape drives where the tape touches the read/write heads, such as 4mm and 8mm, tend to be less reliable (and slower) than tape drives where the tape does not touch the read/write heads, such as 3570 and 3590. If you do implement drives that touch the read/write heads, plan to replace the tapes at regular intervals.

2.5.2 To Library or Not to Library?

ADSM can become tape-intensive. The larger the number and size of clients you are backing up, the greater the load placed on tape. As this load increases, manual tape handling can become time consuming and can create a bottleneck in system performance.

An automated tape library permits an ADSM implementation to run a nearly lights-out operation. Manual intervention should only be required to insert and remove tapes for offsite storage. A tape library having a bar code reader for tape labels reduces errors.

When choosing a tape library, pay attention to the library cycle time, including mount, search, transfer, and dismount time. For busy libraries, a long cycle time can adversely affect performance of ADSM.

We recommend an automated tape library with a bar code reader for all but the smallest of ADSM implementations.

2.5.3 Number of Tape Drives

In attempting to reduce costs, many ADSM customers acquire only one or at most two tape drives. One-drive systems are possible but are subject to outages due to the failure of the sole tape drive. Two-drive systems allow for quicker reclamation and pool copies and reduced drive outages, but restores coming from tape may be held up due to other tape activity. Systems with three or more tape drives can handle restores from tape occurring while tape reclamation or other tape processing is in progress.

A critical consideration for the number of tape drives is how quickly you need to restore one or more clients. More tape drives allow you to restore more clients in a given time. Collocation allows as many clients to be restored simultaneously as you have tape drives. If you have stringent restore requirements, collocation and numerous tape drives make sense.

To calculate the rate of a restore operation, divide the amount of storage to be restored by the sustained data rate of the tape drive (not the instantaneous, or burst, rate) quoted by the manufacturer. If this number does not allow you to meet your service levels, collocation, more tape drives, faster tape drives, a new service level, or another backup strategy may be required.

We recommend a minimum of two tape drives on an ADSM system, with at least three tape drives being preferable.

2.5.4 Number of Tape Volumes

Tape volumes are used to store onsite copies of data, offsite copies of data and database backups. Additional tape volumes are required because all tape volumes are not filled to capacity. Some volumes are required to stock a scratch pool so that mounts for unused tape volumes can be satisfied.

2.5.4.1 Onsite Volumes

To calculate the number of onsite tape volumes required, carry out the following calculations:

1. If this is a sequential storage pool (tape device), multiply the primary storage pool size by the *Number of backup versions* from Table 11 on page 15 to give versions pool size.
2. Add the sum of all *Total storage used* fields for each client from Table 7 on page 9 to the versions pool size, giving tape pool size.
3. Divide the tape pool size by the device capacity to find the number of tape cartridges required.
4. Add 50% to cater for tapes that are in *Filling* status to give the total cartridges required for onsite tapes.
5. If using collocation, there must be at least as many tape cartridges as there are clients.

Example:

1. If the primary storage pool is 1.04 GB, and the number of versions kept is 7, then multiplying $1.04 * 7$ gives a 7.3 GB pool size.
2. If the all the clients are using 0.75, 8, and 25 GB, then $0.75+8+25 +7.3$ equals 41.05 GB.
3. If the tape device has a capacity of 2 GB, then dividing 41.05 by 2 gives 21 cartridges required to store all the data.
4. $21 * 1.5$ gives 32 total cartridges required for onsite tapes.
5. If we use collocation and in this example there are only three clients, we have enough tape volumes.

2.5.4.2 Offsite Volumes

To estimate the number of tape cartridges required for offsite copies, double the number of tape volumes calculated in 2.5.4, "Number of Tape Volumes" on page 31. Keep in mind that collocation may be set on for either, both, or none of onsite and offsite tape pools.

2.5.4.3 Database Tape Volumes

Each database backup requires at least one tape volume. We recommend backing up the database every day and keeping these backups for at least five days.

To calculate the database tape volumes required:

1. Divide the database size calculated in 2.4.2, “ADSM Database” on page 23 by the tape device capacity and round up to the nearest whole number to give the number of tape volumes required for one database backup.
2. Multiply this number by six (five copies plus the copy just being made) to give the total number of tape volumes required for database backups.

For example, If the database size is 350 MB, and the device capacity is 2 GB:

1. $0.350 \text{ GB} / 2 = 0.175 \text{ GB}$ rounds up to 1 tape volume.
2. $1 \text{ tape} * 6 \text{ versions} = 6 \text{ tape volumes}$.

2.5.4.4 Scratch Volumes

A scratch volume is required every time ADSM wishes to write to an unused tape.

To estimate the number of scratch volumes required:

1. Total the number of tape volumes required for onsite tapes, offsite tapes and database backup tapes.
2. Calculate 15% of this number to allocate for scratch tape volumes.

2.5.5 MVS Tape Management

ADSM uses data set names to identify various types of ADSM data sets. The data set name prefix is set by the device class parameter PREFIX. Each device class can have a different data set name prefix. The data set name suffix is fixed by ADSM for various data types. The suffix .DBB indicates a database backup data set. The suffix .BFS indicates an onsite or offsite data copy.

Most tape library management systems on MVS use the data set name to identify tapes to be taken offsite. Since ADSM uses the data set name prefix.BFS for both onsite and offsite copies, the tape management system has no way to identify tapes that must be moved offsite.

To choose another data set name for offsite copies, create another device class for the offsite copies and choose a different prefix. Set the tape library management system to trigger on this different PREFIX.BFS data set name and offsite copies will be identified automatically.

ADSM allows an external data manager (EDM) to control tapes. To inform the EDM when a tape goes scratch, you can use the DFSMSHsm ARCTVEXT parameter. Include the DELETIONEXIT ARCTVEXT parameter in the ADSM server options file. For more information, see *ADSM V3R1 MVS Administrators Guide* (GC35-0277).

If your tape management system uses program names to identify External Data Managers, the ADSM program name is ANRSERV.

2.6 Administrator IDs

Identify who will be the ADSM administrator(s) in your organization.

An ADSM administrator controls ADSM resources. There can be numerous administrators with varying levels of authority. With the advent of the Web backup-archive client, it is now possible to perform backup, restore, archive, and

retrieve operations on the behalf of other users using only a Web browser such as Netscape or Internet Explorer. Help desk personnel can use this to perform these client tasks for their customers without having to logon to the client machine.

If your ADSM installation is large or widely dispersed, you can delegate some authority to administrators based on policy domains or storage pools. Thus a workstation administrator could look after setting data retention criteria for workstation data only (assuming the correct policy domains were set up). A UNIX administrator could be given ADSM authority over UNIX data retention criteria only.

Since ADSM logs all commands issued by administrators and it has no limit on the number of administrators, do not share administrator IDs. Sharing administrator IDs reduces the accountability of each ID, and thus of all the people sharing the ID. Conversely, numerous administrator IDs may give too many people too much authority.

Table 21 on page 33 may help identify the several administrator IDs you want to implement.

Table 21. Administrator IDs Worksheet

Functions	ADSM ID	Authority
Server console	SERVER_CONSOLE	System
System administrator	sysadmin	System
System support	support	System
System reporter	reporter	none
Client administrator	helpdesk	Node

See 7.1, "Administrators" on page 123 for more information about administrators.

2.7 License Considerations

Licenses are used by ADSM in three ways. First, the ADSM server must be licensed to obtain the server code and to legally operate the server. Second, a license is required for each client that ADSM will manage. Third, optional features must be licensed to enable the feature on the ADSM server.

Complete Table 22 on page 33 with the information below for each ADSM server. See 8.1, "Licensed Features" on page 137 for more information.

Table 22. License Requirements Worksheet

	Required
Server type	Windows NT
Network communications	Yes
Hierarchical storage manager	No
Backup-archive clients	2 x 10, 1x 50
Advanced device support	Yes

	Required
Enterprise administration	No
Disaster recovery manager	No
OpenSystems environment	No
Server-to-server virtual volumes	No

2.7.1 Server Type

Enter the ADSM server platform type you are licensing. For example, use Windows NT, AIX, MVS, HP-UX, Sun Solaris, or AS/400.

2.7.2 Network Communications

Enter Yes to license a network communication method like TCP/IP. A network license is not required for the following:

- Administrative client Web interface (HTTP communication method)
- Administrative client access
- Use of a Tivoli or an SNMP receiver

2.7.3 Hierarchical Storage Manager

Enter Yes to license the Hierarchical Storage Manager.

2.7.4 Backup-Archive Client

Enter one or more of the following to add up to the number of ADSM backup-archive clients you wish to administer.

- A x 1 where A is the number of single licenses required.
- B x 5 where B is the number of 5 client licenses required.
- C x 10 where C is the number of 10 client licenses required.
- D x 50 where D is the number of 50 client licenses required.

2.7.5 Advanced Device Support

Enter the number and type of advanced device support licenses required. Look at the URL <http://www.storage.ibm.com/software/adsm/addevice.htm> for more information about what devices are supported.

2.7.6 Enterprise Administration

Enter Yes if multiple ADSM servers will be set up to communicate with each other.

2.7.7 Disaster Recovery Manager

Enter Yes if you wish to use the DRM.

2.7.8 OpenSystems Environment

Enter Yes if you wish to license the OS/390 UNIX Systems Services client.

2.7.9 Server-to-Server Virtual Volumes

Enter Yes if you are setting multiple ADSM servers that will share data with each other.

2.8 Other Considerations

There are numerous other topics to be considered when planning an ADSM installation. Many of these topics are outside the scope of this book, but we are mentioning them for completeness.

- **Staffing:** Staffing requirements need to be addressed. The various functions such as operations, technical support, administration and help desk may all be performed by one person in a small site. Larger sites may find a more specialized approach useful. To provide backup coverage, two people per function is always a good idea.
- **Lead time:** Some tasks such as installing a tape exit in MVS may have considerable lead times before the change can be made. We have highlighted some of these, but check with your technical support group and your change management group for their guidelines.
- **Monitoring:** You may wish to consider monitoring your ADSM server system using a product such as Tivoli. We have highlighted some suggestions for this, but there are many more items that you may wish to monitor. Monitoring also includes monitoring the health of the ADSM software. Numerous queries are useful for displaying information about the ADSM system and its workings. We have included the more basic commands.
- **Charging:** Some ADSM installations charge for their services. This is possible using the accounting records and site specific programs. Some items you may wish to consider charging for include bytes stored, CPU time per client, or tapes used.
- **Code refreshes:** New client code typically has been released every three months. With installations greater than about 50 clients, keeping up with these refreshes of client code becomes a problem. Set up a procedure for tracking which clients are running which release of ADSM code. Design your client installs to be as generic and as similar as possible. If an automated software install process is available, consult with the process owners regarding the best practices to use in setting up ADSM clients.
- **Export/Import:** It is possible to export a client definition and all of its related data from one ADSM server and import it into another ADSM server. This facility is useful for moving clients from one server platform such as Windows NT to another server platform such as AIX. With a large number of clients, or clients with a large amount of data, the export and import can take a significant amount of time, on the order of 24 to 48 hours. In these cases, planning and coordination needs to be done as to when the exporting server ceases backups and the importing server starts backups.
- **Server Scripts:** These are very useful for issuing ADSM commands repeatedly, or with some rudimentary logic around it.
- **SQL queries:** These are powerful queries you can run against the ADSM database to extract information.

- **Problem determination:** Diagnosing and resolving problems are tasks you will have to do on a regular basis. In general, once you have determined you have a problem, install the latest level of ADSM client and ADSM server code and try to recreate the problem. If it still exists, contact your IBM or ADSM representative to search their problem databases. Alternatively, you may post your problem to the ADSM listserv list in hope that someone else has experienced the problem or who may be able to offer some suggestions.
- **Disaster recovery:** We recommend planning and testing for disaster recovery be done on a regular basis. The optional feature Disaster Recovery Manager (DRM) assists in gathering, maintaining and recommending information and planning pertinent to disaster recovery.

Chapter 3. ADSM Server

In this chapter, we explain steps relating to the implementation of an ADSM server. We cover the topics of code installation, options file customization, and other server implementation tasks required for creating the redbook ADSM environment. See 2.3, “Server Architecture Considerations” on page 17 for planning considerations

3.1 Latest Code Updates

ADSM server and client code fixes and enhancements are released on a regular basis. The fixes are available from IBM via the internet or on CD-ROM. If you have good internet service, then we recommend that you download the code. You should be aware that the fixes can be quite large as they are full or near-full code replacements. If you are unable to download from the internet, order the fixes through your usual IBM service channel.

The ADSM Web page, <http://www.storage.ibm.com/software/adsm>, has links to the latest ADSM server fixes, ADSMConnect Agent fixes, ADSM client code, and important download information. Each link leads you to sets of files on the ftp server, index.storsys.ibm.com, describing and containing the fixes. The names and content of these files may vary slightly across platforms:

- `README.ftp`

This file contains the download and install instructions.

- `README.lst`

This file contains important information that is not yet available in manuals. This includes information such as description of code fix, enhancements, limitations, and publication fixes.

- Code image(s)

These files contain the actual code fixes. You should download these files in accordance with the information in the `README.ftp` file.

If you encounter download problems, this Web page links to another page containing download tips that offer possible reasons why you may having problems.

If the fix is unavailable for download, you can use the fix number to order the code from IBM.

We recommend that you keep a copy of the latest server, agent, and client fix files on a suitable file server at your site. This allows easier distribution and code installation, especially for clients. We recommend that for non-MVS systems, you use the ADSM server to store the code.

3.2 Code Installation

You should use the instructions in the associated *Quick Start* manual for your chosen ADSM server platform to install the ADSM code. The *Quick Start* manual is provided with the install media. The latest version of this manual is available from the ADSM Web site in either HTML or PDF format. We recommend that you

download the manual to ensure that you are working with the latest information. The URL for ADSM publications is:

<http://www.storage.ibm.com/software/adsm/pubs/admanual.htm>.

AIX:

Chapter one of the *ADSM V3R1 AIX Quick Start (GC35-0273)* manual contains the steps to install the ADSM code in a section titled *Installing ADSM*. You should install all the components. You should not install the sample command scripts. The sample scripts and our recommended scripts will be installed later.

MVS:

The Program Directory contains the steps to install all the server and administrator client code. The *ADSM V3R1 MVS Quick Start (GC35-0276)* does not contain the steps to install the code but does contain steps for setting up and customizing an ADSM environment.

Windows NT:

Chapter one of the *ADSM V3R1 Windows NT Quick Start (GC35-0295)* manual contains the steps to install all the server and client code. You should accept all the default options.

3.3 Customization

An ADSM Server has a number of options and settings which control its operation. You specify the options in the options file and the settings interactively. ADSM uses the options specified in the options file at server start-up. You specify server settings via an administrative interface.

3.3.1 Options File

ADSM provides a server options file with a set of default options to start the server. You can modify server options by using a text editor. On some platforms, ADSM provides a server options file editor to perform this function. The supplied file contains information on what options and option values can be specified. You can display the current server options via the `query options` command.

We provide recommended options files for the various server platforms, as the layout of the supplied options file is not easy to understand. Appendix B.2, "Server Options Files" on page 249 contains our server option files. The file options are presented in a more logical sequence than those provided during the install. We assume that TCP/IP is the network protocol and that administrators can access the server via a browser such as Netscape or Internet Explorer.

Table 23 on page 39 shows the default location of the ADSM server options file by server platform. We recommend that you take a backup copy of the existing options file before you replace it with your updated version.

Table 23. Server Options File Location

Platform	File Location
AIX	/usr/lpp/adsmserver/bin/dsmserve.opt
HP/UX	/opt/adsmserver/bin/dsmserve.opt
MVS	As specified by DDname OPTIONS
Solaris	/opt/IBMadsm-s/bin/dsmserve.opt
Windows NT	C:\Program Files\IBM\adsm\server\dsmserve.opt

Although the number of server options is very large, there are only a small number that need to be changed for each server. These options fall into the following categories:

- Communication options
- Site-dependent options
- Performance options

Note:

At server initialization, the server reads the server options file. If you update a server option by editing the file, you must remember to stop and restart the server to activate the updated server options file.

3.3.1.1 Communication Options

If your network protocol is not TCP/IP, then we recommend that you update the Communications Options section of our options file to support your protocols. A server can support multiple communication methods. The *Quick Start* manual provides the details on what options and values to specify.

3.3.1.2 Site-Dependent Options

If your locale is not the USA, then we recommend that you update the Date, Number, Time, and Language Options section of the options file to support your requirements. The *Administrative Reference* manual provides the details on what options and values to specify. You must specify the value for each of these options that was determined during the planning phase. Some options do not apply to some server platforms. Refer to the *Administration Reference* manual for specific details. Table 24 on page 39 shows those server options.

Table 24. Server Options: Site Dependent

Option	Description
DATEFORMAT	Specifies the format by which dates are displayed by the server
LANGUAGE	Specifies the language used for messages
NUMBERFORMAT	Specifies the format by which numbers are displayed by the server

Option	Description
TIMEFORMAT	Specifies the format by which times are displayed by the server
DEVCONFIG	The file in which ADSM stores a backup copy of the device configuration information
VOLUMEHISTORY	The file in which ADSM stores a backup copy of the server sequential volume history information
MIRRORWRITE log sequential	Avoids potential log data corruption when operating system fails during log update
ENABLE3590LIBRARY	Support for IBM 3590 tape drives in an IBM 349x automated library

3.3.1.3 Performance Options

Table 25 on page 40 identifies performance options for which we recommend different values than the default. Refer to the *Administration Reference* manual for specific details.

Table 25. Server Options: Performance

Server Option and Value	Description
EXPINTERVAL 0	Specifies the interval, in hours, between automatic inventory expiration runs by ADSM
TXNGROUPMAX 256	Specifies the number of files that are transferred as a group between a client and the server between transaction commit points
MOVEBATCHSIZE 256	Specifies the number of files that are to be moved and grouped together in a batch, within the same server transaction
MOVESIZETHRESH 500	Specifies a threshold for the amount of data moved as a batch, within the same server transaction
LOGPOOLSIZE 2048	Specifies the size of the recovery log buffer pool size
BUFPOOLSIZE 32768	Specifies the size of the database buffer pool
TCPWINDOWSIZE 32	Specifies the size of the buffer used when sending or receiving data.

3.3.2 Settings

ADSM provides default server run-time settings. These settings are stored in the ADSM database and persist across server restarts. You specify the server settings via the administrative interface. You can display the current server settings via the `query status` command.

The default values for the server settings are generally acceptable. There are a few settings for which we recommend values other than the defaults. The settings are changed via the `set` command and fall into the following categories:

- Security related settings
- Our recommended settings

3.3.2.1 Security Related Settings

Table 26 on page 41 shows the various default security related settings for the server. We recommend that you select values for these settings which reflect the security policy at your site.

Table 26. Security Related Settings

Setting and Value	Description
AUTHENTICATION ON	Whether administrators and client nodes must use a password to access the server
INVALIDPWLIMIT 3	Maximum number of failed logon attempts before a node or administrator is locked
MINPWLENGTH 6	Minimum length of a password
PASSEXP 90	Password expiration period
WEBAUTHTIMEOUT 10	Time-out interval for the Web administrative interface

For example, consider the security policy at a company which specifies that all passwords must be changed every 30 days and must be a minimum of 6 characters in length. If three invalid passwords are entered in response to a password prompt, that user must have system access revoked. Administrative users accessing the server via the Web interface are never to be timed out.

The policy is implemented by the following administrative commands:

```
adsm> set authentication on
ANR2095I Authentication parameter set to ON.

adsm> set passexp 90
ANR2092I Password expiration period set to 90 days.

adsm> set minpwlength 6
ANR2138I Minimum password length set to 6.

adsm> set invalidpwlimit 3
ANR2175I Invalid password limit set to 3 attempts.

adsm> set webauthtimeout 0
ANR2198I SET WEBAUTHTIMEOUT: Web authentication time-out set to 0 minutes.
```

3.3.2.2 Recommended Server Settings

We recommend changing some of the default server settings.

The accounting server setting determines whether an accounting record is created every time a client node session ends. The default is not to create these records. We recommend that accounting be switched on to collect these records. For MVS systems, this information is recorded in the Systems Management Facility (SMF). For other systems, this information is written to a file which contains text records that can be viewed directly or can be imported into a spreadsheet program.

All server activity is recorded in an activity log. This log is located in the ADSM database and contains text messages. The log is pruned automatically every night at midnight. The default retention period for the log is only one day. This does not provide an adequate amount of information for production systems. We recommend that you specify a value of between 7 and 14 days.

The administrative command to collect accounting information, keep server activity log records for 10 days, and the results of those commands should look like this:

```
adsm> set accounting on
ANR2091I Accounting parameter set to ON.

adsm> set actlogretention 10
ANR2090I Activity log retention period set to 10 days.
```

3.4 Server Implementation

In this section, we explain how to load the ADSM supplied sample scripts, load the redbook environment scripts, format server volumes, label tape volumes, and use the Windows NT install wizards.

3.4.1 Load Sample Scripts

ADSM is shipped with sample command scripts that can be loaded into the database and run from an ADSM administrative client, administrative Web interface, or server console. They can also be included in ADSM administrative command schedules. The sample scripts, in `scripts.smp`, are primarily SELECT queries, but also include scripts that define volumes, extend the database and recovery log, and backup storage pools.

You can load the sample scripts into the database by using either the DSMSERV LOAD command or the `macro` administrative command. When running the administrative client, it is required that the sample scripts file be accessible from the local machine. The DSMSERV LOAD command can only be used when the server is not running, while the `macro` command can only be used when the server is running. We recommend that you use the `macro` command method.

On non-MVS platforms, the sample scripts file is named `scripts.smp` and is in the install directory for the server. On the MVS platform, the sample scripts file is a member of the ADSM SAMPLIB data set named ANRSCRIPT.

In the following example, we assume that the administrative client is running in a Windows environment and the sample scripts file has been put into the directory `C:\Program Files\IBM\adsm\server`. You must invoke the administrative command line client with the ITEMCOMMIT parameter, otherwise the macro command will fail with a series of error messages. The commands to create the sample scripts and resulting output looks similar to this:

```

C:\Program Files\IBM\adsm\saclient>dsmadm -itemcommit
ADSTAR Distributed Storage Manager
Command Line Administrative Interface - Version 3, Release 1, Level 0.6
(C) Copyright IBM Corporation, 1990, 1997, All Rights Reserved.

Enter your user id: sysadmin

Enter your password: *****

Session established with server PAGOPAGO: Windows NT
Server Version 3, Release 1, Level 2.13
Server date/time: 02-02-1999 17:02:09 Last access: 02-02-1999 17:01:55

adsm> macro C:\Program Files\IBM\adsm\server\scripts.smp
ANS8000I Server command: 'del script q_locked'
ANR1455E DELETE SCRIPT: Command script Q_LOCKED does not exist.
ANS8000I Server command: 'del script q_inv_attempt'
ANR1455E DELETE SCRIPT: Command script Q_INV_ATTEMPT does not exist.

```

3.4.2 Format Server Volumes

The ADSM database, recovery log, and storage pools consist of storage volumes. Each type of volume must be formatted for a particular type of usage. An ADSM server volume is the equivalent of an operating system file.

The process of adding a new server volume consists of two steps:

1. Format the volume for the volume type using the DSMFMT utility.
2. Define the volume to ADSM using the appropriate `define` command.

In a non-MVS environment, those two steps can be combined into a single process: you format the volume and immediately define it in ADSM by using one single storage volume command using the `FORMATSIZE` option. You format and define new database volumes by using the `define dbvolume` command, new recovery log volumes by using the `define logvolume` command, and storage pool volumes by using the `define volume` command. We recommend that you use this method. See sections 4.1, “Defining a New Database Volume” on page 47, 4.2, “Defining a New Recovery Log Volume” on page 49, and 5.8.5, “Defining Storage Pool Volumes” on page 90 for examples.

In an MVS environment, the only way to format server volumes is to use the DSMFMT utility. This utility runs outside of the ADSM server environment. Sample JCL to create and format server volumes is in the ADSM SAMPLIB data set in members named ANRALLO1 and ANRALLO2. The server volume is allocated as a VSAM linear data set. Once the volume has been formatted, it must be defined to ADSM for use.

3.4.2.1 Define Database Volume

For database volumes, the allocated file size is always a multiple of 4MB (plus an additional 1MB for overhead). The command to format a 100MB database volume on a Windows NT system, and the resulting output from that command, should look like this:

```
C:\Program Files\IBM\adsm\utils>dsmfmt -db e:\adsm\database\file07.db 101

ADSTAR Distributed Storage Manager
ADSM for Windows NT Volume Extent/Volume Formatting Program

Licensed Materials - Property of IBM

5639-C59 (C) Copyright IBM Corporation 1990, 1997. All rights reserved.
U.S. Government Users Restricted Rights - Use, duplication or disclosure
restricted by GSA ADP Schedule Contract with IBM Corporation.

Allocated space for e:\adsm\database\file07.db: 105906176 bytes.
```

The command to define the formatted database volume to ADSM, and the resulting output from that command, should look like this:

```
adsm> define dbvolume e:\adsm\database\file07.db
ANR2240I Database volume E:\ADSM\DATABASE\FILE07.DB defined.
```

3.4.2.2 Define Recovery Log Volume

For recovery log volumes, the allocated file size is always a multiple of 4MB (plus an additional 1MB for overhead). The command to format a 24MB recovery log volume on an AIX system, and the resulting output from those commands, should look like this:

```
kindu:[/usr/lpp/adsmerv/bin]$ dsmfmt -log /adsm/recoverylog/file02 25
ADSTAR Distributed Storage Manager/6000
AIX ADSM Server DSMFMT Extent/Volume Formatting Program

Licensed Materials - Property of IBM

5765-C43 (C) Copyright IBM Corporation 1990, 1997. All rights reserved.
U.S. Government Users Restricted Rights - Use, duplication or disclosure
restricted by GSA ADP Schedule Contract with IBM Corporation.

Allocated space for /adsm/recoverylog/file02: 26214400 bytes
```

The command to define the formatted recovery log volume to ADSM, and the resulting output from that command, should look like this:

```
adsm> define logvolume /adsm/recoverylog/file02
ANR2260I Recovery log volume /adsm/recoverylog/file02 defined.
```

3.4.2.3 Define Storage Pool Volume

The command to format a 200MB storage pool volume on an AIX system, and the resulting output from that command, should look like this:

```
kindu:[/usr/lpp/adsm/serve/bin]$ dsmfmt -data /adsm/stgpool/backup1 200
ADSTAR Distributed Storage Manager/6000
AIX ADSM Server DSMFMT Extent/Volume Formatting Program

Licensed Materials - Property of IBM

5765-C43 (C) Copyright IBM Corporation 1990, 1997. All rights reserved.
U.S. Government Users Restricted Rights - Use, duplication or disclosure
restricted by GSA ADP Schedule Contract with IBM Corporation.

Allocated space for /adsm/stgpool/backup1: 209715200 bytes
```

The command to define the formatted storage pool volume in the storage pool named *DISKDATA*, and the resulting output from that command, should look like this:

```
adsm> define volume diskdata /adsm/stgpool/backup1
ANR2206I Volume /adsm/stgpool/backup1 defined in storage pool DISKDATA (device
class DISK).
```

3.4.3 Label Tape Volumes

Tapes used by the ADSM server must have labels written at the start of each volume. Labels are internal records that uniquely identify the tape volume. In an MVS environment, the use of standard labelled tape volumes provides the label information. In non-MVS environments, ADSM must write the labels.

In an MVS environment, any tapes used by ADSM must have MVS standard labels. ADSM does not write those labels; they must be generated by some other method. The remainder of this section does not apply to MVS environments.

In non-MVS environments, you provide the labels on tape volumes by using either the DSMLABEL utility or the `label libvolume` administrative command. The DSMLABEL utility is run outside the ADSM environment, while the `label libvolume` command is run inside the ADSM environment. Both methods work with all tape drives, whether they are stand-alone units or are in libraries. We recommend that you use the administrative command `label libvolume` to label volumes because it writes the labels, does not require removing a drive from ADSM control, and, in the case of a tape library, registers the volume with the server.

See section 12.4.1, “Labelling Tapes” on page 215 for examples of using the `label libvolume` command for both manual and automatic libraries.

3.4.4 Windows NT Configuration Wizards

This section applies to the Windows NT ADSM server only.

The ADSM server utilities are available through the *IBM ADSM* program group. It provides access to various configuration wizards, an administrative interface, a server console, Web sites, and other useful utilities. In particular, the initial configuration wizards provide a structured way to implement an ADSM environment. Each of the configuration wizards is also available independently within the Server Utilities.

Documentation on the use of the initial configuration wizards is in the chapter titled *Performing the Initial Configuration* of the *ADSM V3R1 Windows NT Quick Start* (GC35-0295) manual.

The wizards provide a good front-end to some steps of the ADSM configuration. They are especially useful in the areas of licensing, services configuration, device configuration, volume formatting and media preparation. The configuration wizards are easier to use than the equivalent administrative interfaces. The wizards hide the details of the command interface and provide good help. However, you still need some ADSM knowledge to create a good working environment.

The wizards and server utilities do not cover all aspects of an implementation and give potentially unexpected results in other areas. The wizards do not cover definition of items such as administrative schedules, client option sets, or administrators. The node configuration wizard can easily generate multiple domains, the names of which you have no control in specifying.

The wizards are not suited to bulk entry of definitions. This task is better handled through the use of administrative macros or scripts. A good example of this is the definition of our redbook administrative schedules. The definitions require one administrative command or dozens of window interactions. Our recommendation is to use the Wizards with care.

Chapter 4. Database and Recovery Log

This chapter covers the steps you must run to set up and change your database and recovery log settings to a recommended configuration.

The database and the recovery log are closely related. The ADSM database contains information needed for server operations and information about client data that has been backed up, archived, and space-managed. The database contains pointers to the locations of all client files in the ADSM storage pools. The client data itself is stored in storage pools, not in the database.

We assume that you have planned your database and recovery log sizes. See 2.4, “Disk Considerations and Sizing” on page 23 for planning considerations. If you have not planned the values that your installation requires to run ADSM, we strongly recommend that you do so before continuing.

Use the information you gathered from the planning chapter to issue your own commands. Because the database is highly dependent on site-specific values, we cannot predict all possible combinations. Therefore, we came up with a sequence of commands, so that all you have to do is change the values for those that best fit your environment.

Platform-Specific Commands:

Some minor changes are required in UNIX and Windows because of the different filesystem naming conventions of those platforms (a filesystem in UNIX is a drive in Windows). Because the examples in this section are based on an AIX UNIX machine, if you have a Windows ADSM version, use the Windows directory names instead. For other platforms, see the *ADSM Administrator's Reference* for that server.

4.1 Defining a New Database Volume

ADSM tracks all volumes defined to the database as one logical object. For performance reasons, you can split the database volumes into different disks, so that the actual read and write operations are balanced among the many disks available. For example, Figure 2 on page 48 shows a database that consists of four volumes (`file01` through `file04`). Although the volumes are separate files from an operating system point of view, ADSM tracks the database as a single logical image.

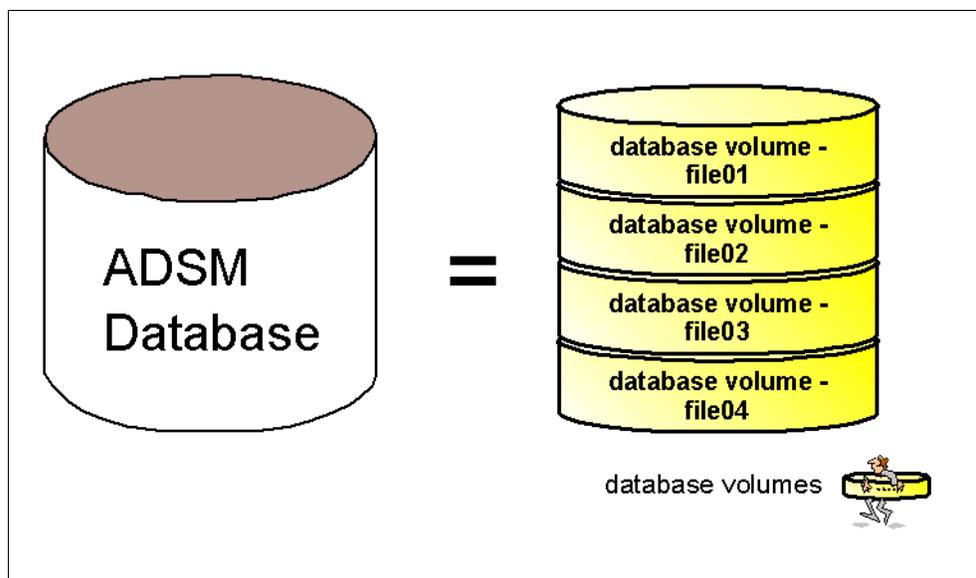


Figure 2. ADSM Database and Its Volumes

All database volumes must reside on random access media. You can have ADSM create the volume before it is assigned.

There are two methods of successfully allocating a new database volume. The most common and easiest to use is to allocate a new ADSM database volume by using the `define dbvolume` command using the `FORMATSIZE` option from either the ADSM console or an administrative client. We recommend that you use this one-step method. The second method is to use the `DSMFMT` utility command to allocate the volume file and then the `define dbvolume` command to define the database volume to ADSM. For further information about the `DSMFMT` utility, see Chapter 3, “ADSM Server” on page 37.

The following command allocates space and defines a new database volume in a single operation. It creates a volume `/adsm/database/primary/file01` of size 176 MB. However, we have to specify 177 M, since the database volume requires 1 MB extra space for overhead.

```

adsm> define dbvolume /adsm/database/primary/file01 formatsize=177 wait=yes

ANR0984I Process 45 for DEFINE DEVOLUME started in the FOREGROUND at 09:13:27.
ANR2240I Database volume /adsm/database/primary/file01 defined.
ANR0986I Process 45 for DEFINE DEVOLUME running in the FOREGROUND processed 1
items for a total of 185,597,952 bytes with a completion state of SUCCESS at
09:13:49.

adsm> query db

Available Assigned Maximum Maximum Page Total Used Pct Max.
Space Capacity Extension Reduction Size Usable Pct Util Pct
(MB) (MB) (MB) (MB) (bytes) Pages Pages Util Util
-----
212 8 204 0 4,096 2,048 1,223 59.7 62.8

adsm> query dbvolume

Volume Name Copy Volume Name Copy Volume Name Copy
(Copy 1) Status (Copy 2) Status (Copy 3) Status
-----
/adsm/database/- Sync'd Undefined Undefined
primary/file01
/usr/lpp/adsmse- Sync'd Undefined Undefined
rv/bin/db.dsm

```

UNIX Users:

The database name is case-sensitive. If you want to create a database volume named `file01`, you must type the name in lower case.

4.2 Defining a New Recovery Log Volume

The recovery log contains information about updates that have not yet been committed to the database. For example, when a client is running backups, all transactional data that controls this operation is first written to the recovery log.

ADSM tracks all volumes defined to the recovery log as one logical object. For performance reasons, you can split the recovery log volumes into different disks, so that the actual read and write operations are balanced to the many disks available. For example, Figure 3 on page 50 shows a recovery log that consists of two volumes (`file01` and `file02`). Although the volumes are separate files from an operating system point of view, ADSM tracks the recovery log as a single logical image.

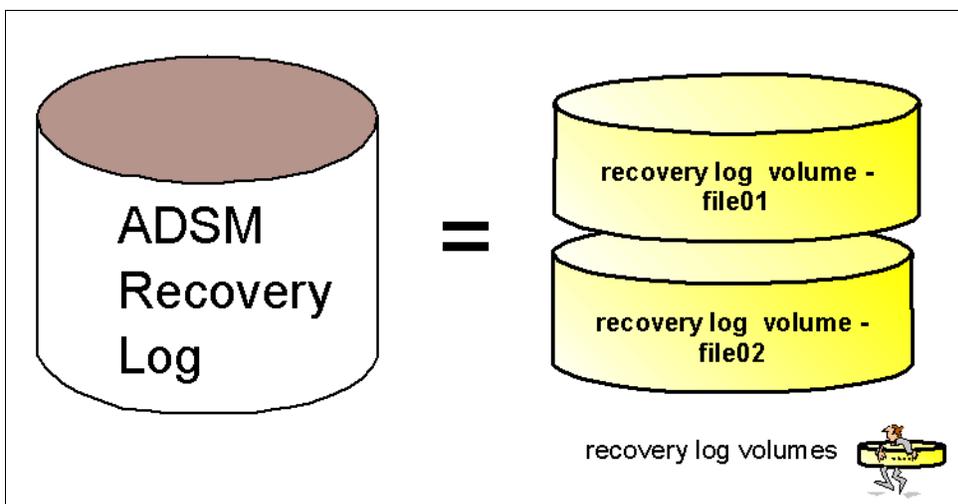


Figure 3. ADSM Recovery Log and Its Volumes

All recovery log volumes must reside on random access media. You can have ADSM create the volume before it is assigned.

There are two methods of successfully allocating a new recovery log volume. The most common and easiest to use is to allocate a new ADSM recovery log volume by using the `define logvolume` command with the `FORMATSIZE` option from either the ADSM console or an administrative client. We recommend that you do this one-step method. The second method is to use the `DSMFMT` utility command to allocate the volume file and then the `define logvolume` command to define the recovery log volume to ADSM. For further information about the `DSMFMT` utility, see Chapter 3, “ADSM Server” on page 37.

The following command allocates space and defines a new recovery log volume in a single operation. It creates a volume `/adsm/log/primary/file01` of size 36 MB. However, we have to specify 37 MB, since the recovery log volume requires 1 MB extra space for overhead.

```

adsm> define logvol /adsm/log/primary/file01 formatsize=37 wait=yes

ANR0984I Process 47 for DEFINE LOGVOLUME started in the FOREGROUND at
10:13:58.
ANR2260I Recovery log volume /adsm/log/primary/file01 defined.
ANR0986I Process 47 for DEFINE LOGVOLUME running in the FOREGROUND processed 1
items for a total of 38,797,312 bytes with a completion state of SUCCESS at
10:14:03.

adsm> query log

Available Assigned Maximum Maximum Page Total Used Pct Max.
Space Capacity Extension Reduction Size Usable Pages Util Pct
(MB) (MB) (MB) (MB) (bytes) Pages -----
-----
44 8 36 4 4,096 1,536 246 16.0 30.7

adsm> query logvol

Volume Name Copy Volume Name Copy Volume Name Copy
(Copy 1) Status (Copy 2) Status (Copy 3) Status
-----
/adsm/log/prima- Sync'd Undefined Undefined
ry/file01
/usr/lpp/adsmse- Sync'd Undefined Undefined
rv/bin/log.dsm

```

UNIX Users:

The recovery log name is case-sensitive. If you want to create a recovery log volume named file01, you must type the name in lower case.

4.3 Setting the Log Mode to ROLLFORWARD

Use the `set logmode` command to set the mode for saving recovery log records. The log mode determines how long ADSM saves records in the recovery log and the kind of database recovery you can use. The two log modes are NORMAL and ROLLFORWARD.

- **NORMAL:** ADSM saves only those records needed to restore the database to the point of the last backup (point-in-time recovery). ADSM deletes any unnecessary records from the recovery log. Changes made to the database since the last backup cannot be recovered. Any backup versions of the database created by issuing the `backup db` command can only be used to perform point-in-time recovery. In NORMAL log mode, you may need less space for the recovery log, because ADSM does not keep all records already committed to the database.
- **ROLLFORWARD:** ADSM saves all recovery log records that contain changes made to the database since the last time it was backed up. ADSM deletes recovery log records only after a successful database backup. The recovery log records can be used to restore a database to its most current state (roll-forward recovery) after loading the most current database backup series. A database backup series created in ROLLFORWARD mode can be used for either point-in-time recovery or roll-forward recovery. We recommended that you enable ROLLFORWARD log mode if your site requires a high level of availability to the ADSM server. ROLLFORWARD log mode may require a significant amount of space to record all activity.

To configure the log mode for ROLLFORWARD, issue the `set logmode` command:

```
adsm> set logmode rollforward
ANR2294I Log mode set to ROLLFORWARD.
```

You can check that the command was successful by using the `query status` command and checking the Log Mode field:

```
adsm> query status
ADSM Server for AIX-RS/6000 - Version 3, Release 1, Level 2.1

Server Name: AIX1
Server host name or IP address:
Server TCP/IP port number:
Server URL:
Crossdefine: Off
Server Password Set: No
Server Installation Date/Time: 02/13/1998 10:59:41
Server Restart Date/Time: 01/22/1999 13:57:58
Authentication: On
Password Expiration Period: 120 Day(s)
Invalid Sign-on Attempt Limit: 3
Minimum Password Length: 0
WEB Admin Authentication Time-out (minutes): 0
Registration: Open
Availability: Enabled
Accounting: On
Activity Log Retention Period: 1 Day(s)
License Audit Period: 30 Day(s)
Last License Audit: 01/23/1999 01:58:01
Server License Compliance: Valid
Central Scheduler: Active
Maximum Sessions: 25
Maximum Scheduled Sessions: 12
Event Record Retention Period: 10 Day(s)
Client Action Duration: 5 Day(s)
Schedule Randomization Percentage: 25
Query Schedule Period: Client's Choice
Maximum Command Retries: Client's Choice
Retry Period: Client's Choice
Scheduling Modes: Any
Log Mode: RollForward
Database Backup Trigger: Not Defined
Active Receivers: CONSOLE ACTLOG SNMP
Configuration manager?: Off
Refresh interval: 60
Last refresh date/time:
```

Note that if you are not using the ADSM database in roll-forward mode presently, but wish to do so, the recovery log pool size will need to be increased. To estimate the new value, reset the cumulative consumption value using the administrative command `reset logconsumption` and then monitor the cumulative consumption over a number of days. Divide the cumulative consumption by the number of days since you reset the value, to get a representative value. A safe size for the log pool should be around 30-40% larger than this figure.

4.4 Defining Database Backup Trigger

Use the `define dbbackuptrigger` command to define settings for the database backup trigger. The database backup trigger determines when ADSM automatically runs a full or incremental backup of the ADSM database and deletes any unnecessary recovery log records.

ADSM uses the settings you specify with this command only when the log mode is set to `ROLLFORWARD`, which you previously configured with the `set logmode` command. With the `define dbbackuptrigger` command, you specify the percentage of the assigned capacity of the recovery log that can be used before ADSM begins a backup of the database. The actual percentage that you must choose is highly dependent on the planning considerations in Chapter 2, “ADSM Implementation Planning” on page 9. We recommend that you use the 75% percentage as a starting point and use the default number of incrementals (six). You must monitor your environment activity to make sure that you do not start unnecessary triggering.

To set the limit of 75% for the recovery log to start a `backup db` and run up to six incremental database backups before a full database backup, issue the `define dbbackuptrigger` command:

```
adsm> define dbbackuptrigger devclass=c3590 logfullpct=75 numincr=6
ANR2282I Database backup trigger defined and enabled.
```

You can check that the command was successful by issuing the `query dbbackuptrigger` command:

```
adsm> query dbbackuptrigger format=detail
      Full Device Class: C3590
      Incremental Device Class: C3590
      Log Full Percentage: 75
      Incrementals Between Fulls: 6
      Last Update by (administrator): ADMIN
      Last Update Date/Time: 11/26/1999 19:13:15
```

4.5 Setting the Expansion Trigger

ADSM lets you fully automate the process of increasing the database and recovery log. For example, assume that you have a 200 MB database and a 100 MB recovery log. You want to increase the size of the database by 25% when 85% is in use, but not to more than 1 GB. You also want to increase the recovery log by 30% when 75% is in use, but not to more than 500 MB.

Note:

Setting a space trigger does not mean that the percentage used in the database and recovery log will always be less than the value specified with the FULLPCT parameter. ADSM checks utilization when database and recovery log activity results in a “commit”. Deleting database volumes and reducing the database does not cause the trigger to activate. Therefore, the utilization percentage can exceed the set value before new volumes are online.

MVS Users:

Database expansion triggering is not fully supported in ADSM MVS. In the MVS environment, the `define spacetrigger` command does not cause new volumes to be created. It causes the database and recovery log to be extended if space is available. Event logging allows messages to display the amount of new space that is needed to satisfy the space trigger utilization parameter. You can use these messages to initiate automatic expansion or to complete an allocation job (refer to ANRFMT1 and ANRFMT, sample jobs that are located in the ASAMPLIB library).

4.5.1 Database Space Trigger

To define a new space trigger for the database in the `/adsm/database/primary` directory (ADSM generates the volume names), issue the `define spacetrigger` command:

```
adsm> define spacetrigger db fullpct=85 spaceexpansion=25 \  
cont> expansionprefix=/adsm/database/primary/ maximumsize=1000  
  
ANR2274I Data Base Space trigger defined and enabled.
```

The MAXIMUMSIZE limit for the database is 999999999 MB. A value of zero, (0) or omitting this parameter disables checking for maximum size. If later on, you need to change the trigger definition, you can use the `update spacetrigger` command.

Whenever the system detects that the database needs extra space, it triggers the expansion as shown in the `query actlog` command:

```
adsm> query actlog  
...  
ANR4414I Data Base Space expansion Is Needed, 6 Megabytes needed.  
ANR4412I Data Base and Recovery Log Space expansion triggered.  
ANR2248I Database assigned capacity has been extended.  
ANR2240I Database volume /adsm/database/primary/D4985547.DBV defined.  
ANR2240I Database volume /adsm/database/primary/D4985547.DBV defined.  
ANR2248I Database assigned capacity has been extended.  
ANR4415I Data Base and Recovery Log Space Expansion Completed.  
...
```

4.5.2 Recovery Log Space Trigger

To define a new space trigger for the recovery log in the `/adsm/log/primary` directory (ADSM generates the volume names), issue the following command:

```
adsm> define spacetrigger log fullpct=75 spaceexpansion=30 \  
cont> expansionprefix=/adsm/log/primary/ maximumsize=500  
  
ANR2279I Recovery Log Space trigger defined and enabled.
```

You can use a `MAXIMUMSIZE` from 9 MB through 5000 MB (5GB) for the recovery log. A value of zero, (0) or omitting this parameter disables checking for maximum size. If later on, you need to change the trigger definition, you can use the `update spacetrigger` command.

Whenever the system detects that the recovery log needs extra space, it triggers the expansion as shown in the `query actlog` command:

```
adsm> query actlog  
...  
ANR4413I Recovery Log Space expansion Is Needed, 4 Megabytes needed.  
ANR4412I Data Base and Recovery Log Space expansion triggered.  
ANR2260I Recovery log volume /adsm/log/primary/L1337547.LOG defined.  
ANR2260I Recovery log volume /adsm/log/primary/L1337547.LOG defined.  
ANR0984I Process 872 for EXTEND LOG started in the BACKGROUND at 09:15:38.  
ANR2268I Recovery log assigned capacity has been extended.  
ANR4415I Data Base and Recovery Log Space Expansion Completed.  
ANR0307I Recovery log extend in progress; 4 megabytes of 4 formatted.  
ANR2268I Recovery log assigned capacity has been extended.  
...
```

4.5.3 When SPACETRIGGER Fails

Although ADSM can expand either the database or the recovery log space if needed, you must pay special attention when mirroring is active, because the `define spacetrigger` command has only one `expansionprefix` option. Thus, if you are running ADSM with software mirrored volumes (DB or LOG), you may get undesirable volumes in a single disk (thus, making mirroring ineffective). To correct this, you must reallocate the mirror copies in the correct place.

Here is an example of database triggering, which leads to an allocation of both image and copy to the same location (`/adsm/database/primary`). In this case, the database volume, `/dsm_db/C0202557.DBV`, must be reallocated to another filesystem, so that mirroring is still adequate:

```
ANR4414I Data Base Space expansion Is Needed, 4 Megabytes needed.  
ANR4412I Data Base and Recovery Log Space expansion triggered.  
NR2240I Database volume /adsm/database/primary/D6102557.DBV defined.  
ANR2241I Database volume copy /adsm/database/primary/C0202557.DBV defined.  
ANR2248I Database assigned capacity has been extended.  
ANR4415I Data Base and Recovery Log Space Expansion Completed.
```

4.6 Mirroring

The following scenario shows the importance of mirroring in the recovery process. As the result of a sudden power outage, a partial page write occurs. The recovery log is now corrupted and not completely readable. Without mirroring, transaction recovery operations cannot complete when the server is restarted. However, if the recovery log is mirrored and a partial write is detected, a mirror volume can be used to construct valid images of the missing pages.

Although you can mirror either the database or the recovery log, we recommend that you mirror both. This gives better availability should you need to recover from a failing disk, as you can see in Figure 4 on page 56.

In this example, the ADSM database volumes are split into two different disks, so if a system failure occurs on one, ADSM can still work. The same thing happens for the recovery log, with each volume duplicated. Note that ADSM still has one database object and one recovery log object (with a mirrored copy).

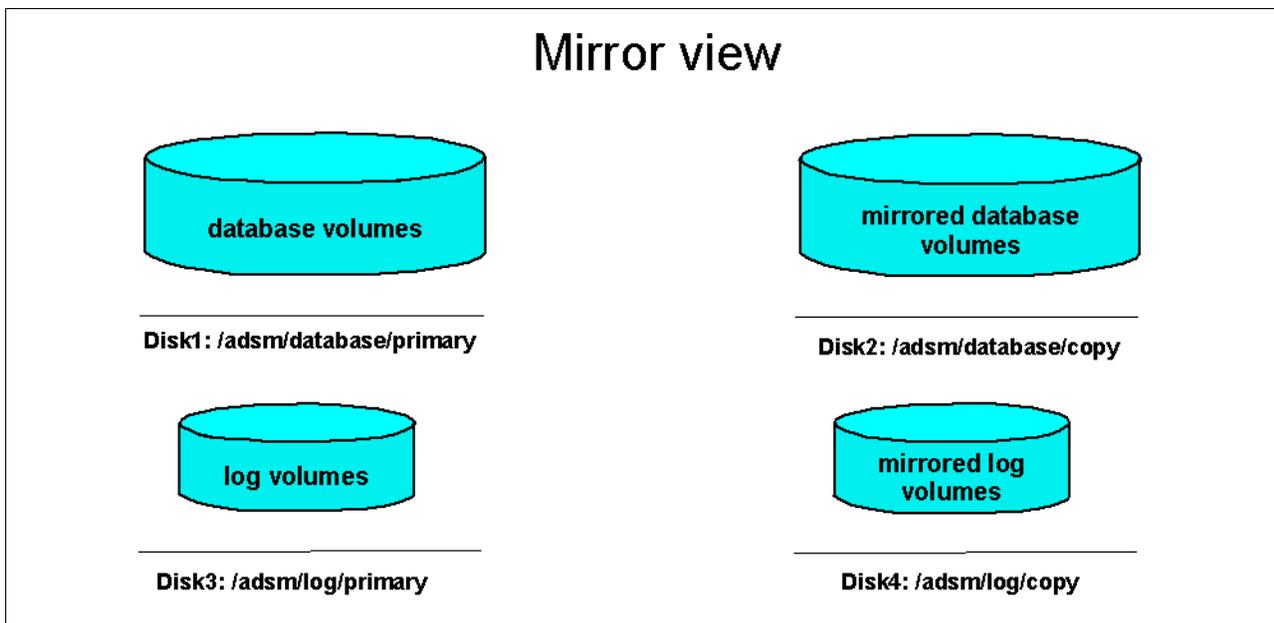


Figure 4. ADSM Mirror View

ADSM mirrored volumes must have at least the same capacity as the original volumes. This means that if your ADSM database is made of two volumes of 100 MB each, then you will need at least two extra 100 MB allocated volumes for the mirrored volumes. If you create volumes larger than necessary, ADSM gives you a warning message (ANR2253W for database and ANR2273W for the log), but it still allows you to use the allocated volume.

You must separate the mirror images so that you do not lose data if the disk fails. We recommend that you use four disks for the database and recovery log volumes to keep each set of volumes on a separate disk. For example, on UNIX, you have all primary database volumes on one disk, `/adsm/database/primary`, and their mirrored volumes on another disk, `/adsm/database/copy`, and the similar distribution for the recovery log volumes and their mirrored copies. Table 27 on page 57 shows a sample database and recovery log allocation for UNIX and Windows respectively:

Table 27. Database and Recovery Log Volumes Allocation

	UNIX location	Windows location
Database volumes	<code>/adsm/database/primary</code>	<code>D:\adsm\database</code>
Mirrored database volumes	<code>/adsm/database/copy</code>	<code>E:\adsm\database</code>
Log volumes	<code>/adsm/log/primary</code>	<code>F:\adsm\log</code>
Mirrored log volumes	<code>/adsm/log/copy</code>	<code>G:\adsm\log</code>

AIX Users:

We recommend that you use ADSM mirroring rather than AIX mirroring. If you use AIX mirroring, you may have a problem with raw volumes, but not with Journaled File System (JFS) files. AIX tracks mirroring activity by writing control information to the first 512 bytes of the USER area in a raw volume. This is not a problem for database and recovery log volumes, but ADSM control information is also written in this area. If AIX overwrites ADSM control information when raw volumes are mirrored, ADSM may not be able to vary the volume online.

4.6.1 Database Mirroring

When you first install ADSM, it does not mirror the database. In our example, the `query dbvolume` command shows the default database (`/usr/lpp/adsmserve/bin/db.dsm`) and one additional volume (`/adsm/database/primary/file01`) created in the previous sections without any mirror images in place:

```
adsm> query dbvolume
```

Volume Name (Copy 1)	Copy Status	Volume Name (Copy 2)	Copy Status	Volume Name (Copy 3)	Copy Status
<code>/adsm/database/primary/file01</code>	Sync'd		Undefined		Undefined
<code>/usr/lpp/adsmserve/bin/db.dsm</code>	Sync'd		Undefined		Undefined

To mirror the database, you must first create all of the new database volumes by using the DSMFMT utility. We create the new copy in `/adsm/database/copy`.

```

kindu:[/usr/lpp/admserv/bin]$ dsfmt -db /adsm/database/copy/file01 201
ADSTAR Distributed Storage Manager/6000
AIX ADSM Server DSMFMT Extent/Volume Formatting Program

Licensed Materials - Property of IBM

5765-C43 (C) Copyright IBM Corporation 1990, 1997. All rights reserved.
U.S. Government Users Restricted Rights - Use, duplication or disclosure
restricted by GSA ADP Schedule Contract with IBM Corporation.

Allocated space for /adsm/database/copy/file01: 210763776 bytes

```

The next step is to define all database volumes in ADSM, using the `define dbcopy` command:

```

adsm> define dbcopy /adsm/database/primary/file01 /adsm/database/copy/file01
ANR2241I Database volume copy /adsm/database/copy/file01 defined.

```

After a database 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.

You can request information about mirrored database volumes by using the `query dbvolume` command. The next screen shows you the database volume synchronized (Sync'd) with its mirrored volume copy:

```

adsm> query dbvolume

```

Volume Name (Copy 1)	Copy Status	Volume Name (Copy 2)	Copy Status	Volume Name (Copy 3)	Copy Status
/adsm/database/primary/file01	Sync'd	/adsm/database/copy/file01	Sync'd		Undefined
/usr/lpp/admserv/bin/db.dsm	Sync'd		Undefined		Undefined

4.6.2 Recovery Log Mirroring

When you first install ADSM, it does not mirror the recovery log. In our example, the command `query logvolume` shows the default recovery log (`/usr/lpp/admserv/bin/log.dsm`) and one additional volume (`/adsm/log/primary/file01`), created in the previous sections, with no mirror in place:

```

adsm> query logvol

```

Volume Name (Copy 1)	Copy Status	Volume Name (Copy 2)	Copy Status	Volume Name (Copy 3)	Copy Status
/adsm/log/primary/file01	Sync'd		Undefined		Undefined
/usr/lpp/admserv/bin/log.dsm	Sync'd		Undefined		Undefined

To mirror the recovery log, you must first create all of the new recovery log volumes by using the DSMFMT utility. Because we want to mirror the /adsm/log/primary/file01 volume, we allocate a new copy for it under /adsm/log/copy/file01:

```

kindu:[/usr/lpp/adsmserve/bin]$ dsmfmt -log /adsm/log/copy/file01 101
ADSTAR Distributed Storage Manager/6000
AIX ADSM Server DSMFMT Extent/Volume Formatting Program

Licensed Materials - Property of IBM

5765-C43 (C) Copyright IBM Corporation 1990, 1997. All rights reserved.
U.S. Government Users Restricted Rights - Use, duplication or disclosure
restricted by GSA ADP Schedule Contract with IBM Corporation.

Allocated space for /adsm/log/copy/file01: 105906176 bytes

```

The next step is to define all recovery log volumes in ADSM, using the `define logcopy` command:

```

adsm> define logcopy /adsm/log/primary/file01 /adsm/log/copy/file01
ANR2261I Recovery log volume copy /adsm/log/copy/file01 defined.

```

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.

You can request information about mirrored recovery log volumes by using the `query logvolume` command. For example, the next screen shows a recovery log with its mirrored images. The recovery log volume is synchronized (Sync'd) with its mirrored volume copy:

```

adsm> query logvolume

```

Volume Name (Copy 1)	Copy Status	Volume Name (Copy 2)	Copy Status	Volume Name (Copy 3)	Copy Status
/adsm/log/primary/file01	Sync'd	/adsm/log/copy/file01	Sync'd		Undefined
/usr/lpp/adsmserve/bin/log.dsm	Sync'd		Undefined		Undefined

Because our recommended configuration removes the default recovery log volume (/usr/lpp/adsmserve/bin/log.dsm), there is no need to mirror it. For further details, see Chapter 4.7.2, “Removing the Default Recovery Log Volume” on page 61.

4.7 Removing Default Volumes

To facilitate the installation procedure, the ADSM database and recovery log volumes are placed in the same location in the disk (`/usr/lpp/adsmserve` on AIX, `C:\Program Files\IBM\adsm\server` on NT). We recommend that you remove the old database and recovery log volumes so that you have a new database and recovery log distribution.

Be sure that you have defined a new database and recovery log volumes as explained in the previous sections before continuing. Otherwise, you can damage your ADSM installation and it will not start.

4.7.1 Removing the Default Database Volume

To remove the default database volume, you must first extend its size by using the `extend db` command, so that there is enough room in the database logical volume. Remember that the ADSM database is a logical object consisting of one or more volumes. Use the `extend db` command to increase the amount of space that can be used by the database within all database volumes previously allocated to ADSM:

```
adsm> extend db 160
ANR2248I Database assigned capacity has been extended.
```

Note:

You can only delete a database volume if there is enough space for the data movement. Therefore, the value in the maximum extension must be at least the size of the volume you want to remove. Otherwise, you will get an error message (ANR2434E).

You can use the `query dbvolume` command to see the full name of the file. You can now delete the default database volume by using the `delete dbvolume` command:

```
adsm> query dbvolume

Volume Name      Copy   Volume Name      Copy   Volume Name      Copy
(Copy 1)         Status (Copy 2)         Status (Copy 3)         Status
-----
/adsm/database/- Sync'd             Undefined         Undefined
primary/file01
/usr/lpp/adsmse- Sync'd             Undefined         Undefined
rv/bin/db.dsm
adsm> delete dbvolume /usr/lpp/adsmserve/bin/db.dsm
ANR2243I Database volume /usr/lpp/adsmserve/bin/db.dsm deleted.
```

The next step is to physically erase the file from the operating system. If you do not do this, you will be keeping unused space in your system. This example assumes use of the default AIX installation directory (`/usr/lpp/adsmserve/bin`). In this case, use the UNIX remove command to delete the unused file:

```
[/usr/lpp/adsmserve/bin]$ rm db.dsm
```

4.7.2 Removing the Default Recovery Log Volume

To remove the default recovery log volume, you must first extend its size by using the `extend log` command, so that there is enough room in the recovery log logical volume. Remember that the ADSM recovery log is a logical object made of one or more volumes. Use the `extend log` command to increase the amount of space that can be used by the recovery log within all recovery log volumes previously allocated to ADSM:

```
adsm> extend log 80

ANR2269I Recovery log extension process initiated (process ID 1201).
ANS8003I Process number 1201 started.
```

Note:

You can only delete a recovery log volume if there is enough space for the data movement. Therefore, the value in the maximum extension must be at least the size of the volume you want to remove. Otherwise, you will get an error message (ANR2445E).

You can use the `query logvolume` command to see the full name for the file. You may now delete the default recovery log volume by using the `delete logvolume` command:

```
adsm> query logvol

Volume Name      Copy   Volume Name      Copy   Volume Name      Copy
(Copy 1)         Status (Copy 2)         Status (Copy 3)         Status
-----
/adsm/log/primary/file01  Sync'd          Undef-          Undef-
                               ined            ined
/usr/lpp/adsmserve/bin/log.dsm  Sync'd          Undef-          Undef-
                               ined            ined

adsm> delete logvolume /usr/lpp/adsmserve/bin/log.dsm

ANR2264I Delete process initiated for recovery log volume
/usr/lpp/adsmserve/bin/log.dsm (process id 1202).
ANS8003I Process number 1202 started.
```

When you ask ADSM to delete a recovery log volume, it moves all valid recovery data that is still on the volume to one of the empty recovery volumes (in our example, `/adsm/log/primary/file01`).

The next step is to physically erase the file from the operating system. If you do not do this, you will be keeping unused space in your system. This example assumes use of the default AIX installation directory (`/usr/lpp/adsmserve/bin`). In this case, use the UNIX `remove` command to delete the unused file:

```
[/usr/lpp/adsmserve/bin]$ rm log.dsm
```

4.8 Database Backup

It is important to run an ADSM database backup. If the database becomes damaged or lost, you can restore it by using the DSMSEV RESTORE DB command to perform ADSM recovery.

If you are following the steps in this chapter, the DBBACKUPTRIGGER will already start database backups as needed. It is a good idea to run a FULL database backup after those database configuration changes we applied in this chapter though, because the trigger may not start in the next few hours.

Use the `backup db` command to back up an ADSM database to sequential access storage volumes. You can use this command to run one of the following types of backup:

- Full backup (TYPE=FULL): Copies the entire ADSM database
- Incremental backup (TYPE=INCREMENTAL): Copies only those database pages that have been added or changed since the last time the database was backed up.

To start a FULL backup:

```
adsm> backup db type=full devclass=8mmdev wait=yes  
  
ANR0984I Process 1200 for DATABASE BACKUP started in the FOREGROUND at  
15:58:05.  
ANR2280I Full database backup started as process 1200.  
ANR4554I Backed up 32 of 959 database pages.  
ANR4554I Backed up 544 of 959 database pages.  
ANR4550I Full database backup (process 1200) complete, 959 pages copied.  
ANR0985I Process 1200 for DATABASE BACKUP running in the FOREGROUND completed  
with completion state SUCCESS at 15:59:35.
```

Note:

The `devclass` parameter may have a different value for your installation. You can run the `query devclass format=detail` command to check the available device classes and `query library` to check the corresponding libraries available for use. You can run the `backup db` command at any time, without having to stop the server.

Note:

We recommend that you send all ADSM database backups to Offsite, so that you have a recovery position for your ADSM environment.

4.9 Additional Commands

Although the recommended settings enable database triggering, you may need to perform some database operations in ADSM. This section shows the possible commands to increase or decrease both database and recovery log space. You

should only run these commands if you experiencing problems with the volume allocations or if you add new volumes to ADSM.

4.9.1 EXTEND DB

Use the `extend db` command to increase the amount of space that can be used by the database within all the database volumes previously allocated to ADSM.

```

adsm>query db
ANR2017I Administrator SERVER_CONSOLE issued command: QUERY DB

Available Assigned Maximum Maximum Page Total Used Pct Max.
Space Capacity Extension Reduction Size Usable Pages Util Pct
(MB) (MB) (MB) (MB) (bytes) Pages Pages Util Pct
-----
20 8 12 0 4,096 2,048 435 21.2 21.2

adsm>extend db 8
ANR2017I Administrator SERVER_CONSOLE issued command: EXTEND DB 8
ANR2248I Database assigned capacity has been extended.
adsm>query db
ANR2017I Administrator SERVER_CONSOLE issued command: QUERY DB

Available Assigned Maximum Maximum Page Total Used Pct Max.
Space Capacity Extension Reduction Size Usable Pages Pages Util Pct
(MB) (MB) (MB) (MB) (bytes) Pages Pages Util Pct
-----
20 16 4 8 4,096 4,096 437 10.7 10.7

```

4.9.2 EXTEND LOG

Use the `extend log` command to increase the amount of space that can be used by the recovery log within all the recovery log volumes previously allocated to ADSM.

```

adsm>query log
ANR2017I Administrator SERVER_CONSOLE issued command: QUERY LOG

Available Assigned Maximum Maximum Page Total Used Pct Max.
Space Capacity Extension Reduction Size Usable Pages Util Pct
(MB) (MB) (MB) (MB) (bytes) Pages Pages Util Pct
-----
108 16 92 12 4,096 3,584 155 4.3 4.7

adsm>extend log 8
ANR2017I Administrator SERVER_CONSOLE issued command: EXTEND LOG 8
ANR0984I Process 12 for EXTEND LOG started in the BACKGROUND at 10:28:55.
ANR2269I Recovery log extension process initiated (process ID 12).
ANR0307I Recovery log extend in progress; 4 megabytes of 8 formatted.
ANR0307I Recovery log extend in progress; 8 megabytes of 8 formatted.
ANR2268I Recovery log assigned capacity has been extended.
ANR0988I Process 12 for EXTEND LOG running in the BACKGROUND processed
8,388,608 bytes with a completion state of SUCCESS at 10:29:04.

ADSM Server for AIX-RS/6000 - Version 3, Release 1, Level 2.1

adsm>query log
ANR2017I Administrator SERVER_CONSOLE issued command: QUERY LOG

Available Assigned Maximum Maximum Page Total Used Pct Max.
Space Capacity Extension Reduction Size Usable Pages Pages Util Pct
(MB) (MB) (MB) (MB) (bytes) Pages Pages Util Pct
-----
108 24 84 20 4,096 5,632 155 2.8 2.8

```

4.9.3 REDUCE DB

Use the `reduce db` command to decrease the amount of space that can be used by the database. To reduce the capacity of the database, you must reduce the database in 4 MB increments. If you do not specify the reduction in 4 MB increments, ADSM rounds the number to the next 4 MB partition. For example, if you reduce the database by 1 MB, ADSM reduces the capacity of the database by 4 MB.

```
adsm> query db
Available Assigned Maximum Maximum Page Total Used Pct Max.
Space Capacity Extension Reduction Size Usable Pages Util Pct
(MB) (MB) (MB) (MB) (bytes) Usable Pages Util Pct
-----
20 20 0 12 4,096 5,120 436 8.5 8.5

adsm> reduce db 8
ANR2250I Database assigned capacity has been reduced.

adsm> query db
Available Assigned Maximum Maximum Page Total Used Pct Max.
Space Capacity Extension Reduction Size Usable Pages Util Pct
(MB) (MB) (MB) (MB) (bytes) Usable Pages Util Pct
-----
20 12 8 4 4,096 3,072 434 14.1 14.1
```

4.9.4 REDUCE LOG

Use the `reduce log` command to decrease the amount of space that can be used by the recovery log. To decrease the capacity of the recovery log, you must reduce the recovery log in 4 MB increments. If you do not specify the reduction in 4 MB increments, ADSM rounds the number to the next 4 MB partition. For example, if you reduce the recovery log by 1 MB, ADSM reduces the capacity of the recovery log by 4 MB.

```
adsm> query log
Available Assigned Maximum Maximum Page Total Used Pct Max.
Space Capacity Extension Reduction Size Usable Pages Util Pct
(MB) (MB) (MB) (MB) (bytes) Usable Pages Util Pct
-----
108 28 80 24 4,096 6,656 150 2.3 2.3

adsm> reduce log 12
ANR2270I Recovery log assigned capacity has been reduced.

adsm> query log
Available Assigned Maximum Maximum Page Total Used Pct Max.
Space Capacity Extension Reduction Size Usable Pages Util Pct
(MB) (MB) (MB) (MB) (bytes) Usable Pages Util Pct
-----
108 16 92 12 4,096 3,584 150 4.2 4.2
```

4.9.5 DELETE VOLHIST TYPE=DBBACKUP

Use the `delete volhistory` command to delete sequential volume history information collected by the server when the information is no longer needed. For example, you may want to delete information about volumes used for obsolete database backups.

When the records that contain volume history information about volumes not in storage pools (volume types BACKUPFULL, BACKUPINCR, and EXPORT) are deleted, the volumes return to scratch status if they were acquired by ADSM as scratch volumes. For scratch volumes with device type FILE, the files are deleted.

When the records that contain volume history information about volumes in storage pools are deleted, the volumes themselves are not affected and remain in the ADSM database.

```
adsm>query volhistory type=dbbackup
ANR2017I Administrator SERVER_CONSOLE issued command: QUERY VOLHISTORY
type=dbbackup

    Date/Time: 01/27/99 15:58:05
    Volume Type: BACKUPFULL
    Backup Series: 12
    Backup Operation: 0
    Volume Seq: 1
    Device Class: 8MMDEV
    Volume Name: DSM001
    Volume Location:

    Date/Time: 02/10/99 10:48:14
    Volume Type: BACKUPFULL
    Backup Series: 13
    Backup Operation: 0
    Volume Seq: 1
    Device Class: C8MM
    Volume Name: ML1802
    Volume Location:

adsm>delete volhistory type=dbbackup todate=today-2
ANR2017I Administrator SERVER_CONSOLE issued command: DELETE VOLHISTORY
type=dbbackup todate=today-2
ANR2467I DELETE VOLHISTORY: 1 sequential volume history entries were
successfully deleted.

adsm>query volhistory type=dbbackup
ANR2017I Administrator SERVER_CONSOLE issued command: QUERY VOLHISTORY
type=dbbackup

    Date/Time: 02/10/99 10:48:14
    Volume Type: BACKUPFULL
    Backup Series: 13
    Backup Operation: 0
    Volume Seq: 1
    Device Class: C8MM
    Volume Name: ML1802
    Volume Location:
```

Chapter 5. Data Storage

ADSM represents data storage with administrator-defined ADSM objects: storage pools and storage pool volumes.

In the following sections, we discuss how ADSM can automatically manage these pools and how you can define them, as well as how to check that the command used has performed the changes you require.

We also briefly discuss some ADSM concepts related to storage pools.

There are example commands you may want to use, including details of the commands needed to follow our recommendations in the checklist.

5.1 Storage Pools

A storage pool is a collection of storage pool volumes, where each storage pool represents one type of media. For example, a storage pool for a 4mm, digital audio tape device (DAT), represents collections of only 4mm tapes, and a storage pool for an IBM 3590 represents collections of only 3590 tapes. A storage pool created on a disk has files formatted under ADSM that are the volumes and are collectively arranged as a storage pool. You can have as many storage pools of the same device type as you need. Figure 5 on page 67 shows the logical relationship between storage pools and their volumes.

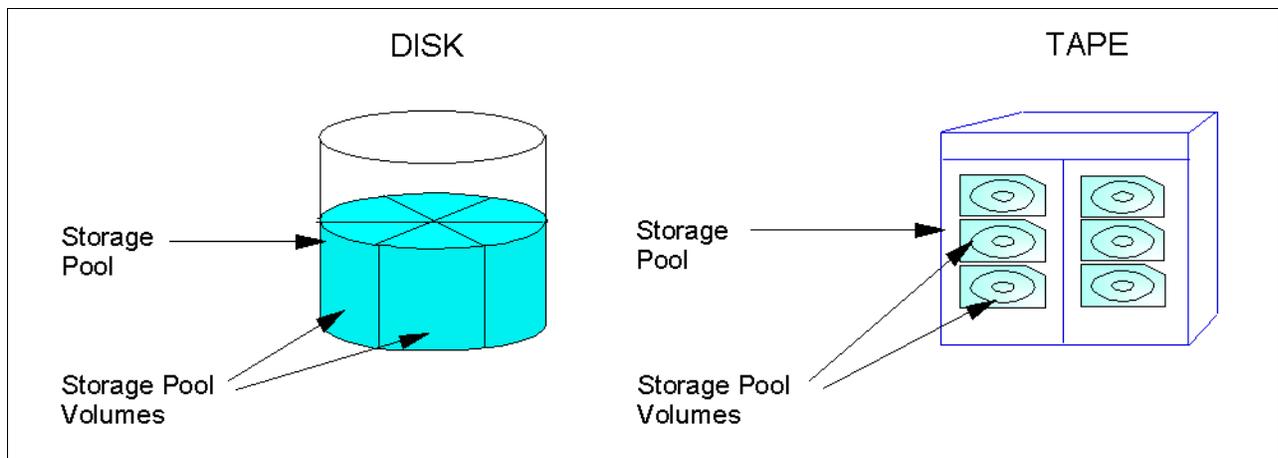


Figure 5. ADSM Storage Pool and Its Volumes

Storage pools for ADSM are defined on a wide range of different devices that are attached locally to the server.

For a complete, up-to-date list of different devices that can be attached to the server platform that you are going to use, the best source is the IBM ADSM Web site. Check the URL <http://www.storage.ibm.com/software/adsm/adsercli.htm>.

Click on the link *ADSM server requirements and ADSM supported storage devices*. Find your ADSM server operating system type and click on the link in the right hand pane to get to our most current device support list.

ADSM has two types of storage pools:

- Primary storage pools
- Copy storage pools

5.1.1 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 or export file data, the requested file is obtained from a primary storage pool if possible. Primary storage pool volumes are always located onsite.

A primary storage pool can use random access storage (DISK device class) or sequential access storage (for example, tape, optical or FILE device classes).

5.1.2 Copy Storage Pools

When an administrator backs up a primary storage pool, the data is stored in a copy storage pool.

The copy storage pool provides a means of recovering from disasters and media failures.

A copy storage pool can use only sequential access storage (for example, tape optical, or FILE device classes).

Copy storage pool volumes can be moved offsite and still be tracked by ADSM. This is done by using the access mode of offsite to ensure that ADSM does not request a volume mount. Moving copy storage volumes offsite provides a means of recovering from an offsite disaster.

The only way to store files in copy storage pools is by using the `backup stgpool` command.

5.2 Storage Pool Hierarchy

There are two main types of devices attached to the ADSM server, random access devices and sequential devices. These two types may be configured in a hierarchy; this allows automation and flexibility to store data on the most suitable media type. You can have an infinite number of levels in a hierarchy, each called a storage pool. You cannot define storage pools pointing to one another in an endless loop.

Copy pools are not part of the storage migration hierarchy. Files are not migrated to or from copy storage pools.

Devices attached to the ADSM server are split into two types: Random access devices and sequential devices.

1. Random access devices refer to magnetic disk devices that are used for two main purposes:
 - To store the database and recovery log.

- To store client data that has been backed up or archived from client nodes. The client data is stored in storage pools.

To protect data in a disk-type storage device we recommend that you use some form of redundant array of independent disks (RAID) either at a hardware or software level. We recommend the use of RAID levels 1, 5, or 0+1.

2. Sequential devices refer to tape devices and/or optical that require manual configuration of the device class and are used for storing backed up and archived client data. ADSM may also store backups of its own database here.

These two main types of devices may be configured in a hierarchy that allows automation and flexibility to store data on the most suitable media type.

A storage pool can be disk, tape, or optical. The management class determines where the client data enters the storage hierarchy. The movement from one storage pool to the next is controlled by the rules that you define in the storage pool. This is illustrated in Figure 6 on page 69.

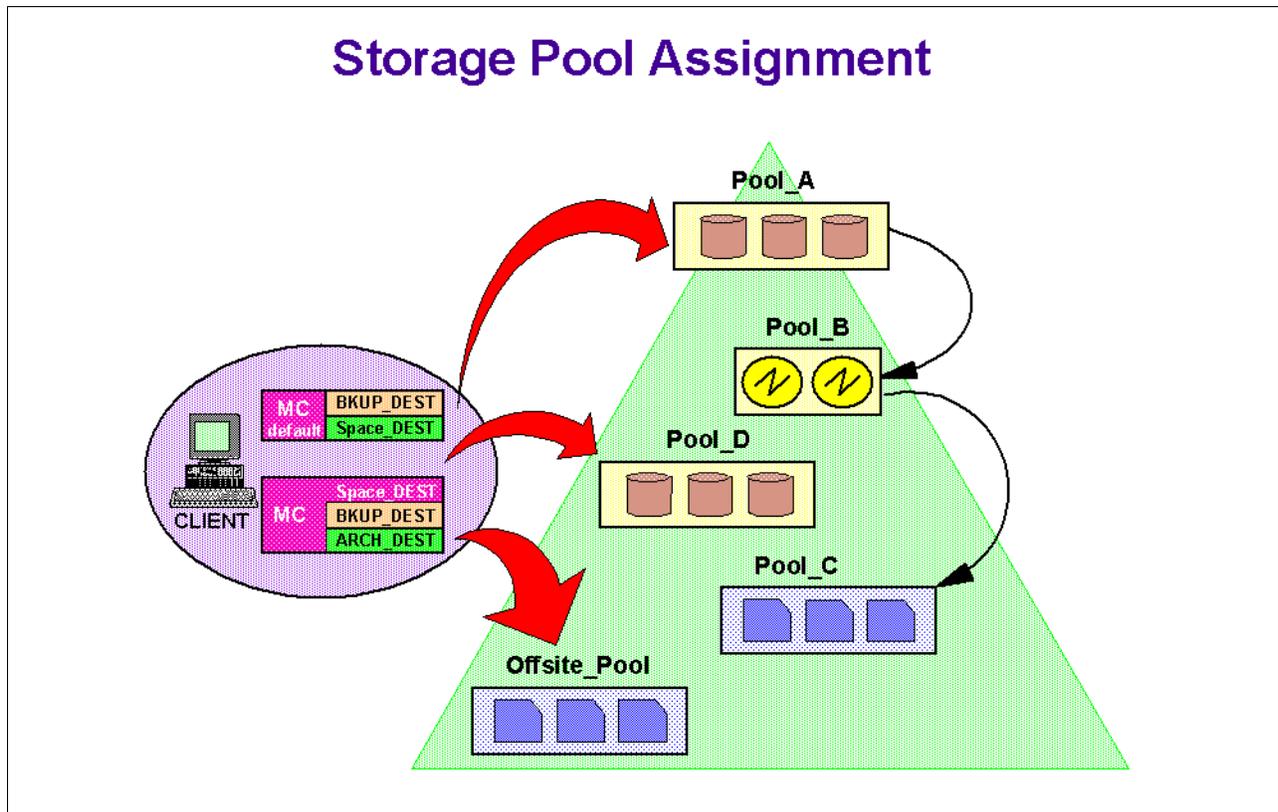


Figure 6. Hierarchical Arrangement of Different Storage Devices

5.3 Movement of Data between Storage Pools

There are two controls available to help you automatically control the space in the storage pools:

- Migration
- MAXSIZE

5.3.1 Migration

Migration helps you control the amount of free space within a storage pool. You can define a high and a low migration threshold for each storage pool in the storage hierarchy. These thresholds tell ADSM when to move data from one storage pool to another.

When the amount of data in a storage pool reaches the high threshold, ADSM moves client data to the next storage pool, until the storage pool reaches the low threshold. This automation mitigates or reduces the chance of the first pool running out of space. The migration process is illustrated in Figure 7 on page 70.

You can specify the number of days to delay migration for files in a storage pool; this ensures that files stay in a pool for a minimum number of days.

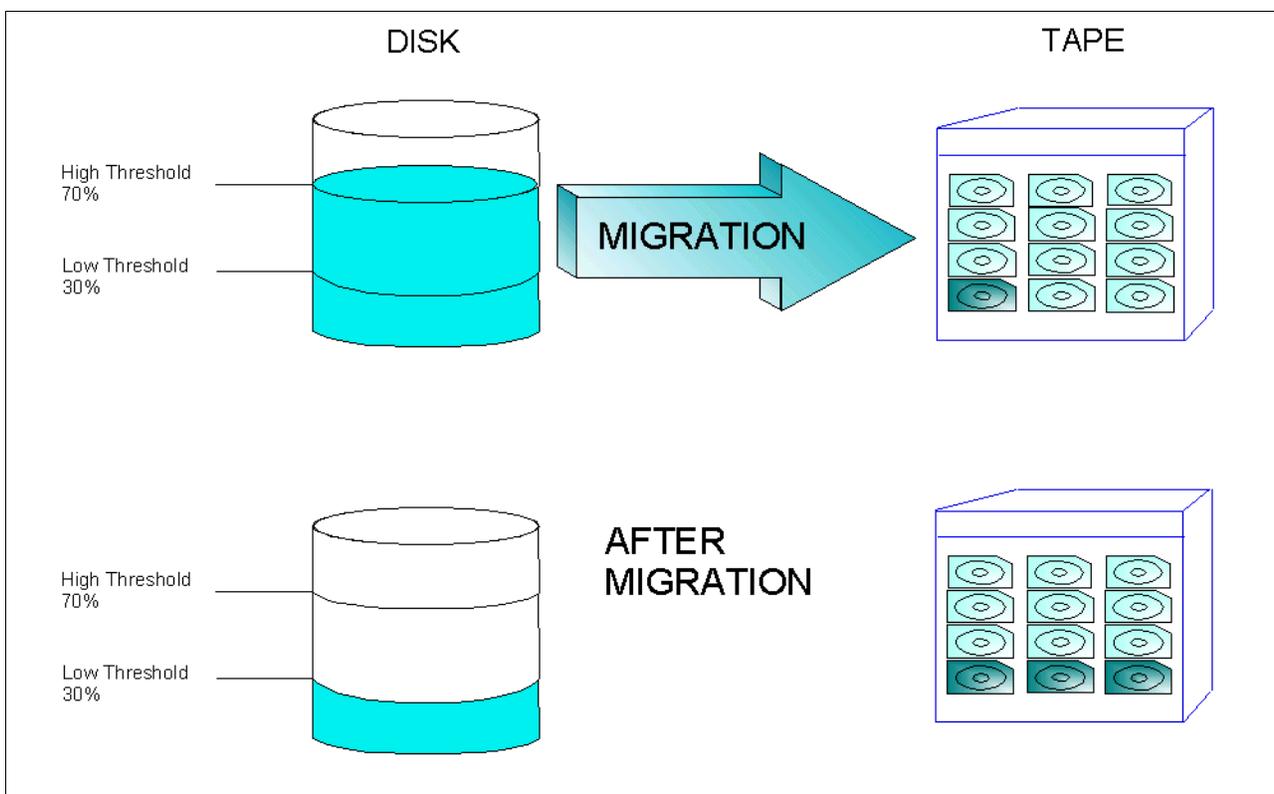


Figure 7. Migration between Storage Pools

5.3.2 MAXSIZE

If you assign a MAXSIZE to a storage pool, files can be placed in that storage pool, provided that they are equal to or less than the assigned MAXSIZE. If a file is too large for the first storage pool, it goes straight to the next storage pool defined.

MAXSIZE is not required. If you do not assign a MAXSIZE, and a file entering the storage pool causes it to exceed its high threshold, migration occurs to move data to the next storage pool.

5.4 Reclamation

Reclamation is used for sequential storage pools to free complete tape volumes. Because ADSM keeps a defined number of versions of files as it does incremental backups, the oldest copy of a file, beyond the defined number of versions to keep, gets marked for expiry.

This file will then be deleted when ADSM next performs expiration processing. It is common that a tape volume will have files that will expire on different dates. Therefore as these files reach their expiry date and the expiration process occurs, empty spaces will appear on the tape volume.

The amount of space that can be reclaimed on a volume increases as files on the volume are deleted. When the percentage of space that can be reclaimed on a volume rises above the value for the RECLAIM option, the volume is eligible for reclamation. Active files on the volume are rewritten to other volumes in the storage pool, making the original volume available for new files.

We recommend that you specify a value of 50 percent or greater for this parameter so that files stored on two volumes can be combined onto a single output volume.

See Figure 8 on page 71 for an illustration of the tape reclamation process.

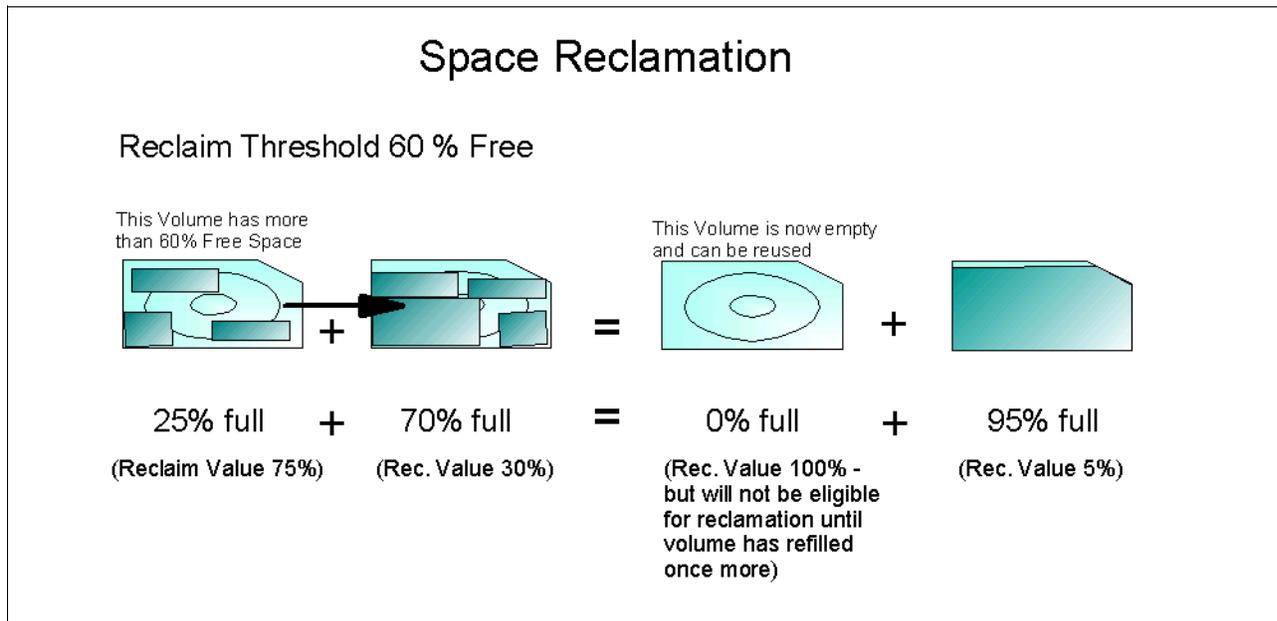


Figure 8. Space Reclamation of Tape Volumes

5.4.1 Single Drive Reclamation

Reclamation can be performed by a single drive in ADSM Version 3 by specifying the RECLAIMSTGPOOL parameter. This parameter points to another storage pool that can be used as the holding area for the data being consolidated.

The storage pool specified as the reclaim storage pool can be any primary storage pool on the system or a new primary storage pool created for this

purpose. The only disk pool allowed is the one whose DEVTYPE parameter is set to FILE.

When the amount of reclaimable space on a media volume exceeds the reclamation threshold, ADSM automatically begins the reclamation process. The volume to be reclaimed is mounted in the drive, and the active data is moved to the reclaim storage pool. If the reclaim storage pool is filled, the volume being reclaimed is dismounted, and a new volume in the same tape pool is mounted. The reclaimed data in the reclaim pool is then migrated to that tape volume. Once this process is complete, it repeats, until all valid data has been reclaimed from the source volume being reclaimed.

When defining the reclaim storage pool, you must also define the NEXTSTGPOOL parameter pointing back to the pool being reclaimed. Thus the reclaimed data can be migrated back to the original storage pool.

5.4.2 Reclamation of Offsite Volumes

Care must be taken if reclamation is set for volumes that are offsite.

If the volumes are actually offsite, in a copy pool, then ADSM cannot physically move the data from one volume to another, as they are in a vault, not a tape drive. The way it manages to perform reclamation for an offsite copy pool is by obtaining the active files from a primary storage pool or from an onsite volume of a copy pool. These files are then written to a new volume in the copy pool and the database updated. A message is then issued that an offsite volume was reclaimed. The new volume will get moved to the offsite location, and the offsite volumes, whose active data it has combined, will be moved back to the scratch pool onsite, when the re-use delay parameter has been satisfied.

However, if a disaster occurs, we could have a problem. If the reclaimed volumes have already been returned to the scratch pool, and the new volume has not yet been taken offsite, then we have lost any offsite backup. This is the reason for the *delay-reuse* period for tapes returning to the scratch pool.

This is achievable by controlling when reclamation will occur by scheduling it properly. Basically, you must ensure that you have performed a backup on all the onsite storage pools to their offsite copy pools, then perform a database backup so the offsite database copy points to all the files in their new locations. Then you can perform reclamation. The delay-reuse period must be long enough to ensure that the reclaimed volumes do not return to the scratch pool before another database backup is taken.

5.5 Reduce Restore Times

There are two methods of reducing restore times that we discuss here.

- Collocation: This restores data from the minimum number of tape volumes.
- Disk caching: This maintains data on disk after it has been migrated.

5.5.1 Collocation

Collocation is a process by which the server attempts to keep all files belonging to a client node on a minimal number of sequential access storage volumes. This process is illustrated in Figure 9 on page 73.

To have ADSM collocate data when files from different client nodes are mixed in the same storage pool, set COLLOCATE 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 or retrieve many files from the storage pool. Collocation thus improves access time for these operations. However, collocation will increase the number of mounts and number of tape volumes used during backup.

You can choose to have the ADSM server collocate at the client or filesystem level.

When collocation is disabled, the server attempts to use all available space on each volume before selecting a new volume. Although this process provides better utilization of individual volumes, user files can become scattered across many volumes; therefore, restoring a client completely may require many volume mounts.

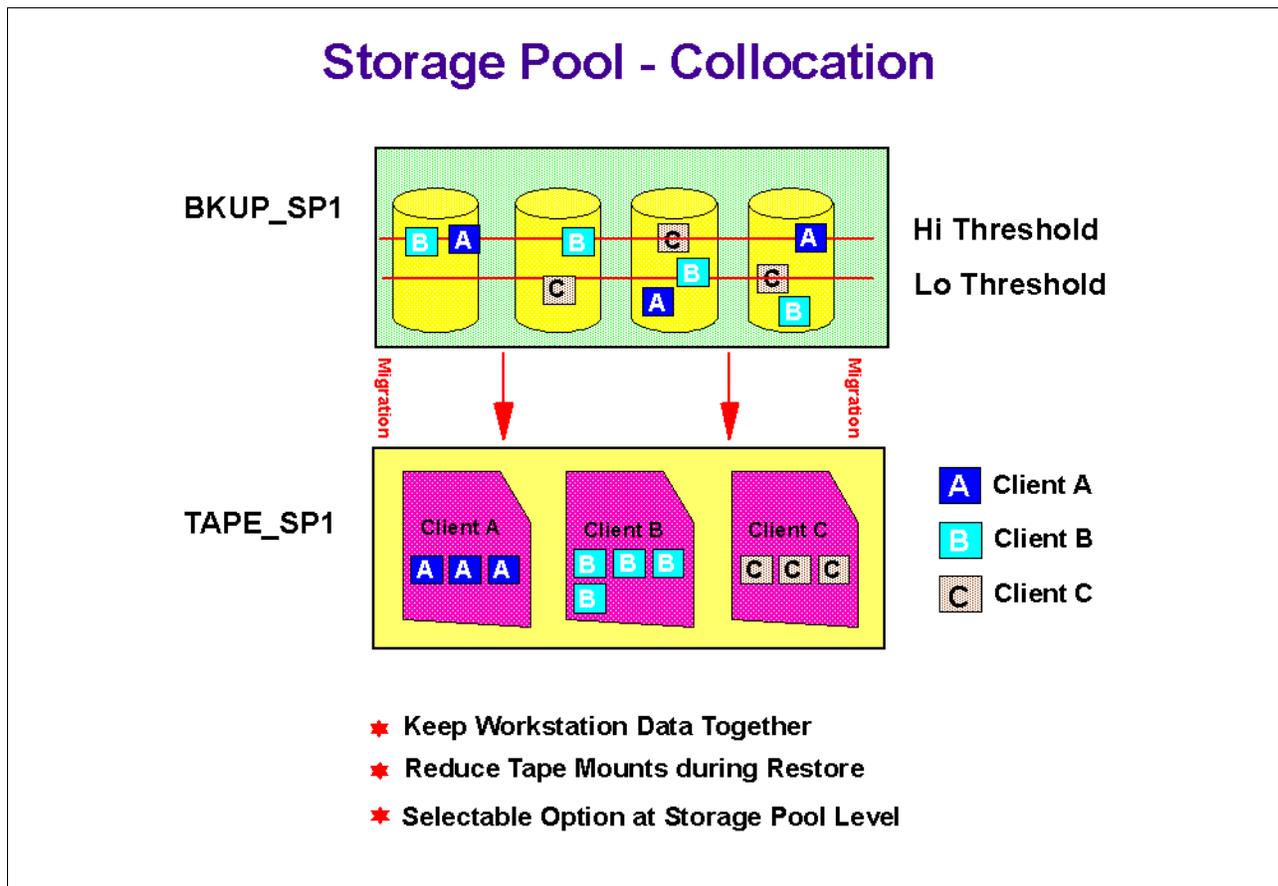


Figure 9. Storage Pool Collocation on a Client Level

5.5.2 Disk Caching

Disk caching is used for disk storage pools only. When cache is enabled (CACHE=YES), the migration process leaves behind duplicate copies of the files on disk after the server migrates these files to subordinate storage pools in the storage pool 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. When space is needed, the server reclaims space by writing over the cached files. Files that have the oldest retrieval date and occupy the largest amount of disk space are overwritten first.

Do not use cache if you have limited database space. When you use cache, more database space is needed because the server has to keep track of both the cached copy of the file and the new copy in the subordinate storage pool. Cache may also slow backup times and be of little value if your available storage pool space is smaller than the amount of changed data to be backed up. This is because the disk pool will be full of cached copies of data to begin with, and so the backup operation will have to create space for every write to this pool. The backup operation will cause a migrate anyway, and will then fill again. The cached data left after a typical end of day migration will only represent a portion of the amount of daily changed data.

5.6 RAID

To protect data in a disk type storage device, we recommend that you use some form of redundant array of independent disks (RAID).

RAID can be implemented at either a hardware or software level. Protecting data stored on a disk is a subject in itself, and below we have just touched on some of the possibilities you may want to look into further. The best solution for you will depend on many factors.

A hardware RAID solution can be beneficial to reduce the performance penalty on the server, when implementing RAID. It can also provide disks and controllers that can be exchanged on the fly, should they fail. This solution can also protect against failed controllers, power supply units, and the input power supply itself, and it can provide a spare disk to automatically take over, should a disk fail. You may also consider implementing disk storage that is shared among a number of servers with the ability to allocate and increase storage for each server, as the need arises, without any server downtime.

A software type RAID solution can also provide levels of redundancy with a much lower initial outlay if implemented carefully, using a number of physical disks.

There are also different ways of transmitting data to the physical disks, such as SCSI, SSA and fiber, which all have advantages and disadvantages.

Whichever path you choose, we recommend that you implement one of the three following RAID levels:

- RAID 1 Mirroring
- RAID 1 + 0 Mirror and Stripe
- RAID 5 Distributed Parity

RAID can only provide protection from hardware failure. Software failure is still possible.

5.7 Setting Up Storage Pools

To define a storage pool initially, you need to judge how large it needs to be. See 2.4, “Disk Considerations and Sizing” on page 23 and 2.5, “Tape Considerations and Sizing” on page 29 for planning considerations. If your initial calculations are inaccurate or data to be backed up grows over time, the storage pool can be increased in size easily and dynamically.

When defining a new device to the ADSM server, the device must first be configured under the operating system. Once we have a device name the operating system understands, we can refer to this, when configuring the device under ADSM. Physical devices are treated under ADSM as two types: Random Access Devices and Sequential devices, and these are connected locally to the server.

- Random access devices

Random access devices refer to magnetic disk devices that are used for two main purposes:

- To store the database and recovery log.
- To store client data that has been backed up or archived from client nodes. The client data is stored in disk storage pools.

ADSM stores data on magnetic disk in two ways:

- Random access volumes
- Storing data sequentially using a FILE device class.

The physical part of random access devices do not need defining to ADSM; they can be referred to directly using a drive/path reference. The device type of file used is also predefined in ADSM.

- Sequential storage devices

Sequential storage devices under ADSM refer to tape libraries, optical devices and write once read many (WORM) devices

The drive and library must first be defined under the operating system. In the following example on AIX, we define an ADIC VLSDLT library, the drive is set at SCSI ID 6, the library is set at SCSI ID 3. First you have to find out the address of the SCSI card by issuing the following command:

```
kindu:[/]$ lsdev -Cc tape
rmt1 Available 00-00-0s-6,0 Other SCSI TApe Drive
```

Now that you know the address of the SCSI card 00-00-0s and the ID and logical unit number (LUN) of the drive 6,0 you can define the drive and library to AIX using the ADSM driver.

To define the drive, start the system management interface tool (SMIT) and go to Devices -> ADSM Devices -> Tape Drive -> Add a Tape Drive. Pick the `ADSM-SCSI-MT` device listed. Pick the card with address 00-00-0s. Supply the connection address in the brackets [6,0]

To define the library, in SMIT go to Devices -> ADSM Devices -> Library -> Add a Library. Pick the `ADSM-SCSI-LB` device listed. Pick the card with the address 00-00-0s. Supply the connection address in the brackets [3,0].

Windows NT Users:

On Windows NT, you define drives and libraries under Settings -> Control Panel -> Tape Devices -> Drivers.

To configure the device under ADSM, five components must be configured. These can be thought of in two component classes, physical and logical. These components must be configured in the order of the five steps following.

- Physical
 1. The library: AUTOMATED or MANUAL.
 2. The drive (or drives, if there is more than one in your library)
- Logical
 3. The device class: Here you specify the device type and associate the device class name with the device type and the library name. Here the device type is from an ADSM defined list, and the command is common to all ADSM server platforms, for example 3570,4mm,or DLT.
 4. The storage pool: This specifies the device class to use for the storage pool you are creating and includes specific storage pool rules for handling data.
 5. The volumes: This is where we label the volumes and either define them as scratch volumes or as private volumes belonging to a particular storage pool.

To define storage pools to ADSM successfully, there is an order that must be followed. If you try to create an ADSM object that has a parameter referring to another item, that item must exist, and therefore it should be created first.

Table 28 on page 76 details which steps you have to complete in order to configure each component for the device and storage pool you wish to set up.

Table 28. Guide to Setting Up Storage Pools

Order of Definition	Component Class	Component	To Define the Type of Device Below, Follow Each Component Section				
			Manual Library	Automated Library	Disk Storage Pool	Sequential Disk Storage Pool	Copy Storage Pool
1	Physical	Library	6.7.1	6.8.1	N/A	N/A	N/A
2		Drive	6.7.2	6.8.2	N/A	N/A	N/A
3	Logical	Device Class	6.7.3.1	6.8.3	6.7.5.2.	6.7.5.3	6.7.5.4
4		Storage Pool	6.7.4.1	6.8.4	6.7.6.2	6.7.6.3	6.7.6.4
5		Volume	6.7.5.1	6.8.5	6.7.7.2	N/A	6.7.7.3

5.7.1 Defining a Library — Physical

In this section we refer to a manual library. See section 6.8.1 for an automated library.

First of all, you need to define a library, even if you are using a single tape drive. The following command defines a manual library named M8MM. There is no device number associated with a manual library because the library part can be thought of as virtual.

```
adsm> define library m8mm libtype=manual
ANR8400I Library M8MM defined.

adsm> query library m8mm

Library Name: M8MM
Library Type: MANUAL
Device:
Private Category:
Scratch Category:
External Manager:
RSM Media Type:
```

The private and scratch category's apply to 3494/3495 type libraries only. The external manager is the path to the location of the external library manager. If you have one, this is where some external media management software identifies the appropriate drive for media access operations. You would need to specify LIBTYPE=EXTERNAL for this parameter to be required.

5.7.2 Defining a Drive — Physical

In this section we refer to a manual library. See section 6.8.2 for an automated library.

After defining the library, you should now define a drive for the library to use.

The drive to use for the manual library must be specified, using its device number that the operating system recognizes it by. The drive therefore must be configured under the operating system.

Windows NT Users:

Run the ADSM Server Utilities, and under Device Information, check the ADSM Device Name, for example mt5.0.0.2. Depending on the device type, you need to install additional device drivers.

AIX Users:

Run the following AIX command to find the device name, for example /dev/mt5:

```
kindu:[/]$ lsdev -C -c tape
```

Depending on the device type, you need to install Atape support.

The following command defines a drive named DRV5 in the manual library on Windows NT named M8MM, the tape drive is at SCSI ID 5, LUN 0, port 0, bus 2.

```

adsm> define drive m8mm drv5 device=mt5.0.0.2
ANR8404I Drive DRV5 defined in library M8MM.

adsm> query drive

```

Library Name	Drive Name	Device Type	Device	ON LINE
M8MM	DRV5	8MM	MT5.0.0.2	Yes

Having defined the physical parts of the library to ADSM, we can now define the logical part to complete the process. There are three steps to this:

- Device Class
- Storage Pool
- Volumes

5.7.3 Defining a Device Class — Logical

A device class can be thought of as a software driver. It is defined so ADSM can communicate with the physical drive correctly.

5.7.3.1 Manual Library

To find the different device types available, look in the administrators reference manual or use the online help facility by issuing the `help define devclass` command. This is the online help facility available in ADSM to give detailed information on every parameter of a command.

The following command defines a device class named C8MM that uses a predefined ADSM device type of 8MM in the library named M8MM. When a tape is sitting idle in a drive in the library, it remains there for 5 minutes (MOUNTRETENTION=5). It uses a tape format of 8500 that gives a capacity of 5GB on a 112m tape.

```

adsm> define devclass c8mm devtype=8mm library=m8mm format=8500 mountretent=5
ANR2203I Device class C8MM defined.

adsm> query devclass

```

Device Class Name	Device Access Strategy	Storage Pool Count	Device Type	Format	Est/Max Capacity (MB)	Mount Limit
C8MM	Sequential	0	8MM	8500	0.0	DRIVES
DISK	Random	4				

The best place to look for information on tape format parameters is in the administrators reference manual. A more common setting is `FORMAT=DRIVE`, which allows ADSM to select the highest format that can be supported by the drive on which a volume is mounted.

5.7.3.2 Disk Storage Pool

We do not need to define the device class for a disk storage pool, because ADSM has a predefined device class of DISK.

5.7.3.3 Sequential Disk Storage Pool

This is where a disk can simulate a sequential device, such as tape.

If you are using a single drive and want to perform reclamation, then the pool type you specify as the reclaim storage pool must be a primary sequential storage pool. This is one instance where you may want to use a sequential disk storage pool with a device type of FILE.

A copy pool must also use sequential access storage and can be defined on a disk using this device type.

To configure a sequential storage pool, you first need to define a device class of FILE, and tell ADSM which directory to use and what the maximum size of this file can be.

The following command defines a device class named CFILE using an ADSM predefined device type of FILE with a maximum capacity of 12MB.

```
adsm> define devclass cfile devtype=file directory=/dsm_stg3 maxcapacity=12M
ANR2203I Device class CFILE defined.
```

```
adsm> query devclass cfile
```

Device Class Name	Device Access Strategy	Storage Pool Count	Device Type	Format	Est/Max Capacity (MB)	Mount Limit
CFILE	Sequential	0	FILE		12.0	1

5.7.3.4 Copy Storage Pool

A copy storage pool can use only sequential access storage (for example, tape, optical, or file device classes). The device class you should define has to represent the type of device you are using for your copy pool. If you are using the same library as one already defined, then you can use the same device class that has been defined for it — we could have used CDLT in our example. If you are using a separate library or an MVS server platform, you need to define a separate device class.

A separate library needs a separate device class to identify it.

MVS Users:

Most tape library management systems on MVS use the dataset name to identify tapes to be taken offsite. ADSM uses the format `<prefix>.BFS` for both onsite and offsite dataset names in one device class. To choose another dataset name for offsite copies, create another device class for the offsite copies and choose a different prefix. Set the tape library management system to trigger on this offsite name, and offsite copies will be identified automatically.

The following command defines a device class named COFFSITE for library LDLT using an ADSM predefined device type of DLT. In this case, this is the type of library we use to produce our offsite copies. The device class we define, COFFSITE, is used when we define the copy storage pool itself.

```

adsm> define devclass coffsite devtype=dlr format=drive library=ldlt
ANR2203I Device class COFFSITE defined.

adsm> query devclass

```

Device Class Name	Device Access Strategy	Storage Pool Count	Device Type	Format	Est/Max Capacity (MB)	Mount Limit
DLT	Sequential	0	DLT	DRIVE	0.0	DRIVES
COFFSITE	Sequential	0	DLT	DRIVE	0.0	DRIVES
DISK	Random	3				

5.7.4 Defining a Storage Pool — Logical

ADSM has two types of storage pools:

- Primary storage pools
- Copy storage pools

A primary storage pool can use random access storage (DISK device class) or sequential access storage (for example, tape, optical or file device classes). A copy storage pool can use only sequential access storage.

5.7.4.1 Primary Storage Pool for a Manual Library

Following the device class, you can now define the storage pool you are going to use with the library.

The following command defines a storage pool named TAPEDATA using the device class named C8MM. The high migration parameter is set to 100 to stop migration because there is no NEXTSTGPOOL for this storage pool to migrate to. The MAXSCRATCH value of 10 means that when 10 scratch tapes have been used, the storage pool indicates that it is full — this may not be desirable. However, this parameter is used to estimate the total number of volumes in the pool and the corresponding estimated capacity of the pool. MAXSCRATCH is a required value for sequential type storage pools.

```

adsm> define stgpool tapedata c8mm highmig=100 maxscratch=10
ANR2200I Storage pool TAPEDATA defined (device class C8MM).

adsm> query stgpool tapedata

```

Storage Pool Name	Device Class Name	Estimated Capacity (MB)	Pct Util	Pct Migr	High Mig Pct	Low Mig Pct	Next Storage Pool
TAPEDATA	C8MM	0.0	0.0	0.0	100	70	

5.7.4.2 Primary Storage Pool for Disk

Using the ADSM predefined device class of DISK, you can define a primary disk storage pool.

The following command defines a storage pool named DISKDATA with a high migration threshold of 70% and a low migration threshold of 30%. It uses the predefined ADSM device class of DISK.

```
adsm> define stgpool diskdata disk description="Data Storage" high=70 low=30
ANR2200I Storage pool DISKDATA defined (device class DISK).
```

```
adsm> query stgpool diskdata
```

Storage Pool Name	Device Class Name	Estimated Capacity (MB)	Pct Util	Pct Migr	High Mig Pct	Low Mig Pct	Next Storage Pool
DISKDATA	DISK	0.0	0.0	0.0	70	30	

5.7.4.3 Primary Storage Pool for Sequential Disk

This is where a disk can simulate a sequential device, such as tape.

The following command defines a storage pool named DISKFILE using a device class named CFILE. Because it simulates a sequential storage pool, the MAXSCRATCH parameter must be included, and as there is no next storage pool in this case, migration is disabled.

```
adsm> define stgpool diskfile cfile maxscratch=100 highmig=100
ANR2200I Storage pool DISKFILE defined (device class CFILE).
```

```
adsm> query stgpool diskfile
```

Storage Pool Name	Device Class Name	Estimated Capacity (MB)	Pct Util	Pct Migr	High Mig Pct	Low Mig Pct	Next Storage Pool
DISKFILE	CFILE	0.0	0.0	0.0	100	70	

Difference in Operation

The server opens a file when one is required, and so the storage pool shows 0.0 Estimated Capacity until data has been sent to it. Following a client backup operation to this storage pool, we can see the capacity increase. Checking the volume used for this operation, the server has created a file named 00000083 with an extension of .BFS to store a client's data. Because of this difference in operation, volumes do not have to be defined for this type of diskpool — the server creates files itself, within the overall maximum capacity of the device class. Note: There are suffixes other than .BFS that may be created when using a sequential pool for other types of operation.

```
adsm> query stgpool diskfile
```

Storage Pool Name	Device Class Name	Estimated Capacity (MB)	Pct Util	Pct Migr	High Mig Pct	Low Mig Pct	Next Storage Pool
DISKFILE	FILE	1,200.0	0.6	1.0	100	70	TAPEDATA

```
adsm> query volume stgpool=diskfile
```

Volume Name	Storage Pool Name	Device Class Name	Estimated Capacity (MB)	Pct Util	Volume Status
/dsm_stg3/00000083.BFS	DISKFILE	FILE	12.0	59.7	Filling

Single Drive Reclamation Using Disk for Reclaim Storage Pool

Reclamation can be performed by a single drive in ADSM version 3 by specifying the RECLAIMSTGPOOL parameter. This parameter points to another storage pool that can be used as the holding area for the data being consolidated. The holding area then refers back to the original pool being reclaimed as its next pool to migrate to. See Chapter 6.4.1 for more information about reclamation.

The following command updates the existing storage pool named TAPEDATA to use the disk storage pool named DISKFILE to hold data temporarily when performing reclamation.

```
adsm> update stgpool tapedata reclaimstgpool=diskfile  
ANR2202I Storage pool TAPEDATA updated.
```

```
adsm> query stgpool tapedata format=detailed
```

```
Storage Pool Name: TAPEDATA  
Storage Pool Type: Primary  
Device Class Name: CDLT  
Estimated Capacity (MB): 40,960.0  
Pct Util: 1.6  
Reclaim Storage Pool: DISKFILE
```

The disk storage pool with a device class of FILE, named DISKFILE, is then updated to point back to the pool being reclaimed.

```
adsm> update stgpool diskfile nextstgpool=tapedata  
ANR2202I Storage pool DISKFILE updated.
```

```
adsm> query stgpool diskfile
```

Storage Pool Name	Device Class Name	Estimated Capacity (MB)	Pct Util	Pct Migr	High Mig Pct	Low Mig Pct	Next Storage Pool
DISKFILE	FILE	0.0	0.0	0.0	100	70	TAPEDATA

5.7.4.4 Copy Storage Pool

This pool has a reuse delay of five days before reclaimed offsite tape volumes are moved back to the scratch pool. This is to allow for a long company shutdown, such as Easter holidays, where the offsite data must remain intact in case a disaster occurs onsite and it is required. You must set this value equal to the longest period of time before normal operations resume for you. Normal operations (database and primary pools) are backed up and volumes are taken offsite.

The following command defines a copy storage pool named OFFDIRS with device class COFFSITE. Reclamation is switched off and the reuse delay is set for five days.

```
adsm> define stgpool offdirs coffsite pooltype=copy reclaim=100 \  
cont> maxscratch=10000 reusedelay=5  
ANR2200I Storage pool OFFDIRS defined (device class COFFSITE).  
  
adsm> query stgpool pooltype=copy
```

Storage Pool Name	Device Class Name	Estimated Capacity (MB)	Pct Util	Pct Migr	High Mig Pct	Low Mig Pct	Next Storage Pool
OFFDIRS	COFFSITE	0.0	0.0				

5.7.5 Defining Storage Pool Volumes - Logical

ADSM is dynamic; you can add or remove volumes without interrupting server operations. For example, if you install a new type of device as a disk pool, the ADSM administrator can move the storage pool data from the old pool to the new pool without shutting down the server. Or, if you have to add space to a storage pool, you can easily define new volumes and thereby expand the size of the storage pool without disrupting service.

Library volumes are volumes checked in to an automated library, including scratch volumes and private volumes.

5.7.5.1 Manual Library

There are two methods of labelling a volume in a manual library. Described here is the more efficient one-step process, using the `label libvolume` command. The second method uses the `DSMLABEL` utility to create a volume and the `define volume` command from within ADSM to define the volume to a storage pool.

Mount the volume you want to label in the library. The following command labels a tape in the library named M8MM with a label V00001 and overwrites any existing label on the tape.

```
adsm> label libvolume m8mm v00001 overwrite=yes  
ANS8003I Process number 3 started.  
  
adsm> query process
```

Process Number	Process Description	Status
3	LABEL LIBVOLUME	ANR8804I Labelling volume V00001 in library M8MM.

5.7.5.2 Disk Storage Pool

Disk volumes can be defined to an ADSM storage pool using either a one-step or two-step process. Described here is the more efficient one-step process. The two-step process requires preparing the disk volume using the DSMFMT utility and then defining the volume to the storage pool using the `define volume` command.

The following command defines a volume that is located in `E:\ibm\adsm\stgpools` and that is named `diskdata1.dsm` to the storage pool DISKDATA. It also prepares the volume and formats it as a background process of size 50 MB.

```
adsm> define volume diskdata e:\ibm\adsm\stgpools\diskdata1.dsm formatsize=50
ANR2491I Volume Creation Process starting for
e:\ibm\adsm\stgpools\diskdata1.dsm, Process Id 9.
```

This results in our storage pool, DISKDATA, looking like this:

```
adsm> query stgpool diskdata
Storage      Device      Estimated   Pct   Pct   High   Low   Next
Pool Name    Class Name  Capacity   Util  Migr  Mig    Mig  Storage
              (MB)
-----
DISKDATA     DISK        50.0      0.0   0.0   70    30
-----

adsm> query volume stgpool=diskdata

Volume Name          Storage      Device      Estimated   Pct   Volume
Pool Name            Class Name  Capacity   Util  Status
              (MB)
-----
E:\IBM\ADSM\STGPOOLS\DI- DISKDATA     DISK        50.0      0.0   On-Line
SKDATA1.DSM
```

Defining additional volumes to the storage pool is how you increase the size of the storage pool dynamically with no interruption to the ADSM service.

Note:

If we had defined a "Next Storage Pool" to migrate to, this storage pool would have to be defined first.

5.7.5.3 Copy Storage Pool

If you are not using a disk for your copy storage pool, then you need to define volumes for your copy storage pool. This is done by using the methods described for defining volumes in a manual library or an automatic library (5.7, "Setting Up Storage Pools" on page 75).

5.8 Recommended Setup

Chapter 2, “ADSM Implementation Planning” on page 9 , may have referred you to this section of the data storage chapter. To complete this section successfully, the component parts of your ADSM setup must be followed, in order, from 5.8.1, “Defining an Automated Library” on page 85 through 5.8.5, “Defining Storage Pool Volumes” on page 90.

5.8.1 Defining an Automated Library

Automated libraries require a device number to be associated with them. This is the ID of the robot that moves the tapes within the library.

The library must be specified, using its device number that the operating system recognizes it by. The library therefore must be configured under the operating system.

Windows NT Users:

Run the ADSM Server Utilities. Under Device Information, check the ADSM Device Name, for example lb0.0.0.0.

AIX Users:

Run the following AIX command to find the device number, for example, /dev/lb0:

```
kindu:[/]$ lsdev -C -c library -H
```

The following command defines our recommended solution, an automated library, of library type SCSI , named LDLT, that was seen by our AIX system at /dev/lb0.

```
adsm> define library ldlt libtype=SCSI device=/dev/lb0
ANR8400I Library LDLT defined.
```

```
adsm> query library ldlt
```

Library Name	Library Type	Device	Private Category	Scratch Category	External Manager
LDLT	SCSI	/dev/lb0			

5.8.2 Defining a Drive in an Automated Library

Libraries with multiple physical drives must have each drive defined to ADSM and also the element address specified. The element address gives a unique name to each part of the library that concerns ADSM. There are element addresses for each physical drive, the robot, and each cartridge slot within the library.

5.8.2.1 Element Information

To find element information look at the IBM Web site:

<http://www.storage.ibm.com/software/adsm/adsercli.htm#servreq>

Click on the supported device list for your operating system and click on the library you are attaching. The device support worksheet, shown in Figure 10 on page 86, is a useful document to complete and file as part of your installation.

Device Support Worksheet

This device support worksheet can be used to record SCSI IDs and device names for the devices that you are attaching to your ADSM server system. Where needed, the worksheet also shows the element numbers (addresses) for drives, slots, and robotics in libraries.

ADIC VLS DLT

Device	SCSI ID	Device Name
Tape drive 1 (element 2)	_____	/dev/mt_____
Autochanger (element 0)	_____	/dev/lb_____

Figure 10. Device Support Worksheet from ADSM Web Site

5.8.2.2 Drive

The following command defines a drive named DRVL in library named LDLT. The drive is at /dev/mt1 and has an element address of 2. For libraries with a single drive, the element address is an optional parameter. It is shown in the example for completeness only.

```

adsm> define drive ldlt drv1 device=/dev/mt1 element=2
ANR8404I Drive DRV1 defined in library LDLT.

adsm> query drive ldlt drv1 format=detailed

      Library Name: LDLT
      Drive Name: DRV1
      Device Type: DLT
      Device: /dev/mt1
      ON LINE: Yes
      Element: 2
Last Update by (administrator): ADMIN
Last Update Date/Time: 02/04/99 16:31:03
Cleaning Frequency (GB): NONE

```

5.8.3 Defining a Device Class to an Automated Library

For our recommended solution we have set up two device classes here. One is called CDLT for the DLT library we are using in our example, and an optional one is called COFFSITE for the offsite storage pools.

Note:

The library LDLT must have two drives for these two device classes to work.

The following command defines a device class named CDLT for library LDLT using an ADSM predefined device type of DLT. A tape sitting idly in one of the drives in the library remains there for 15 minutes before being automatically dismounted.

```

adsm> define devclass cdltd devtype=dlt format=drive library=ldlt mountret=15
ANR2203I Device class CDLT defined.

adsm> query devclass

```

Device Class Name	Device Access Strategy	Storage Pool Count	Device Type	Format	Est/Max Capacity (MB)	Mount Limit
CDLT	Sequential	0	DLT	DRIVE	0.0	DRIVES
DISK	Random	3				

Carefully consider the value you use for mount retention. If your setup requires frequent mounts of different volumes, using an automatic library, then set a low value. A value of 0 may even be beneficial for drives in constant use. Collocation increases the number of tape mounts required during backup. At the other end of the scale, with a manual drive and only one volume in daily use, console messages requesting a volume to be mounted are unwanted; therefore a high value should be used. The value range is 0-9999 minutes. Other factors that may influence your choice are: the time it takes to mount a volume on the library you are using, and regular access to the same volume. Keeping the volume in the drive may improve response time here.

We now set up the optional separate device class for offsite storage pools. The device class you should define has to represent the type of device you are using for your copy pool. If you are using the same library as one already defined, then you can use the same device class that has been defined for it — we could have used CDLT in our example. If you are using a separate library or an MVS server platform, you need to define a separate device class.

Most tape library management systems on MVS use the data set name to identify tapes to be taken offsite. ADSM uses the format <prefix>.BFS for both onsite and offsite data set names in one device class. To choose another data set name for offsite copies, create another device class for the offsite copies and choose a different prefix. Set the tape library management system to trigger on this offsite name and offsite copies will be identified automatically.

The following command defines a device class named COFFSITE for library LDLT using an ADSM predefined device type of DLT.

```

adsm> define devclass coffsite devtype=dlt format=drive library=ldlt
ANR2203I Device class COFFSITE defined.

adsm> query devclass

```

Device Class Name	Device Access Strategy	Storage Pool Count	Device Type	Format	Est/Max Capacity (MB)	Mount Limit
CDLT	Sequential	0	DLT	DRIVE	0.0	DRIVES
COFFSITE	Sequential	0	DLT	DRIVE	0.0	DRIVES
DISK	Random	3				

5.8.4 Defining Storage Pools

There are four primary storage pools and two copy pools to set up in our recommended solution. Client data backups go to storage pool DISKDATA and are then migrated to storage pool TAPEDATA before going offsite to copy pool OFFDATA. Client directory structures, which are much smaller in size, go to storage pool DISKDIRS and are then backed up to the offsite copy pool OFFDIRS. See Figure 11 on page 88 for our recommended storage pool setup.

In 1.2, “Redbook Support Material” on page 2, we show how to load a predefined macro into ADSM. The macro `mac.stgcreate` that we provide to create the storage pools in our redbook environment, is shown in Appendix B.1.6, “Create Storage Pools” on page 248.

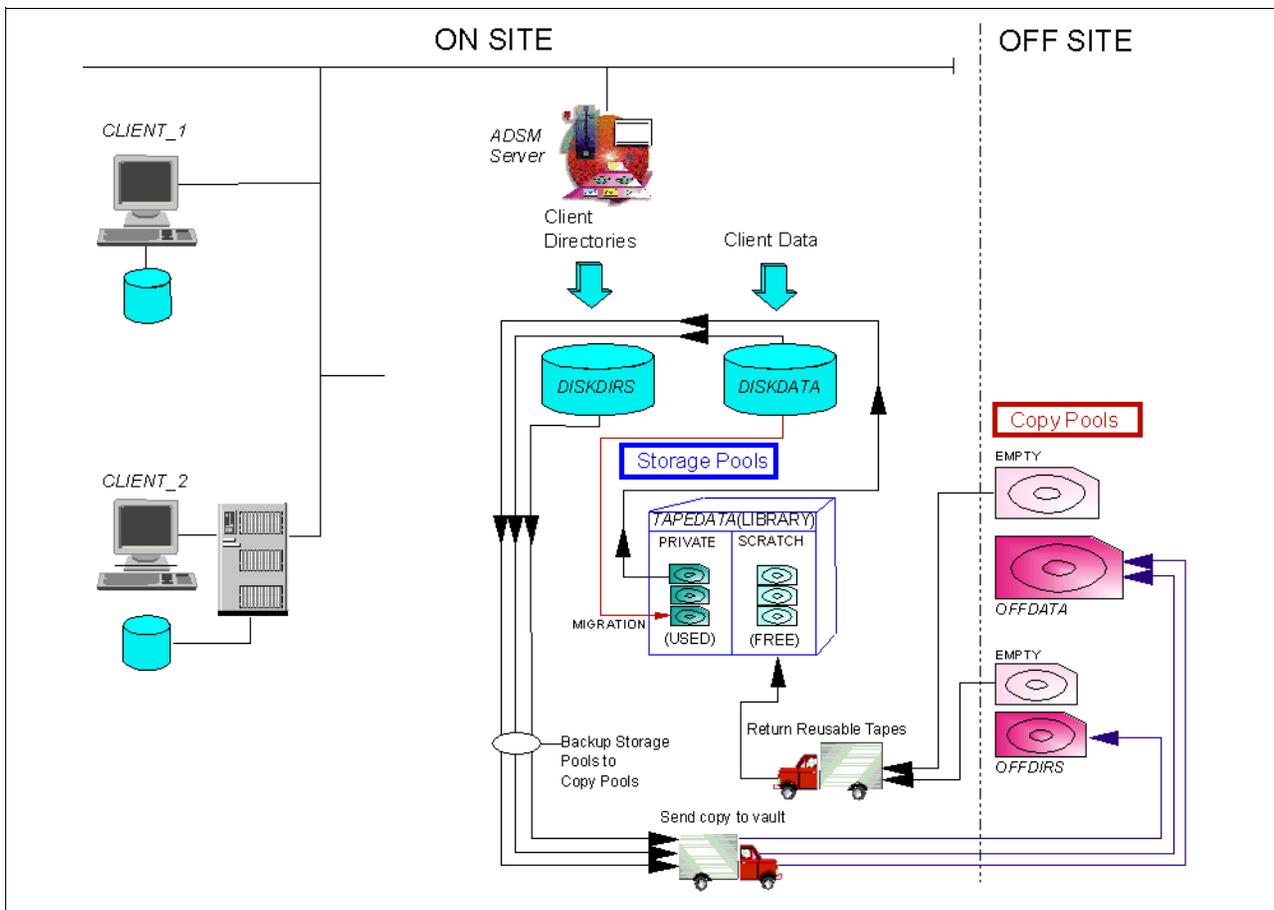


Figure 11. Storage Pool Hierarchy for Our Recommended Setup

5.8.4.1 Primary Storage Pools

For our recommended solution we set up four primary storage pools: TAPEDATA, DISKDATA, DISKDIRS, and NONE. Storage pool TAPEDATA must be set up before storage pool DISKDATA. This is because DISKDATA refers to TAPEDATA as its next storage pool.

The following commands set up these storage pools:

1. A primary storage pool named TAPEDATA, with device class CDLT. It has collocation and reclamation turned off, and has a delay of one day before empty tapes are reused. To define a storage pool for an automated library, use the same procedure as defining a pool for a manual library.
2. A primary storage pool named DISKDAT, with a device class of DISK. It has an upper threshold of 70% that must be reached for migration to start, and a lower threshold of 30% that is reached or surpassed before migration stops. The pool to which data migrated goes is called TAPEDATA. The disk pool is not using disk caching.
3. A primary storage pool named DISKDIRS, with a device class of DISK. It has migration turned off.
4. A dummy storage pool named NONE, which is just an empty definition. No data will be stored in this storage pool. There are no storage pool volumes assigned to it. See 6.2.3, “Defining Management Classes” on page 115 for further information.

```

adsm> define stgpool tapedata cdlt highmig=100 maxscratch=10000 collocate=no \
cont> reclaim=100 reusedelay=1
ANR2200I Storage pool TAPEDATA defined (device class CDLT).

adsm> define stgpool diskdata disk nextstgpool=tapedata highmig=70 lowmig=30 \
cont> cache=no
ANR2200I Storage pool DISKDATA defined (device class DISK)

adsm> define stgpool diskdirs disk highmig=100
ANR2200I Storage pool DISKDIRS defined (device class DISK)

adsm> define stgpool none disk
ANR2200I Storage pool NONE defined (device class DISK).

adsm> query stgpool

```

Storage Pool Name	Device Class Name	Estimated Capacity (MB)	Pct Util	Pct Migr	High Mig Pct	Low Mig Pct	Next Storage Pool
DISKDATA	DISK	0.0	0.0	0.0	70	30	TAPEDATA
DISKDIRS	DISK	0.0	0.0	0.0	100	70	
NONE	DISK	0.0	0.0	0.0	90	70	
TAPEDATA	CDLT	0.0	0.0	0.0	100	70	

5.8.4.2 Copy Storage Pools

For our recommended solution we have set up two copy storage pools: OFFDIRS and OFFDATA. These pools have a reuse delay of five days before reclaimed offsite tape volumes are moved back to the scratch pool. This is to allow for a long company shutdown, such as Easter holidays, where the offsite data must remain intact in case a disaster occurs onsite and it is required for the recovery. You must set this value equal to the longest period of time before normal operations resume for you. Normal operations (database and primary pools) are backed up and volumes are taken offsite.

The following commands set up these storage pools:

1. A copy storage pool named OFFDIRS, with device class COFFSITE. Reclamation is switched off and the reuse delay is set for five days
2. A copy storage pool named OFFDATA, with device class COFFSITE. Reclamation is switched off and the reuse delay is set for five days. The MAXSCRATCH value is set high, purposefully, to avoid a "storage pool full" message occurring.

```
adsm> define stgpool offdirs coffsite pooltype=copy reclaim=100 \  
cont> maxscratch=10000 reusedelay=5  
ANR2200I Storage pool OFFDIRS defined (device class COFFSITE).  
  
adsm> define stgpool offdata coffsite pooltype=copy reclaim=100 \  
cont> maxscratch=10000 reusedelay=5  
ANR2200I Storage pool OFFDATA defined (device class COFFSITE).  
  
adsm> query stgpool pooltype=copy
```

Storage Pool Name	Device Class Name	Estimated Capacity (MB)	Pct Util	Pct Migr	High Mig Pct	Low Mig Pct	Next Storage Pool
OFFDATA	COFFSITE	0.0	0.0				
OFFDIRS	COFFSITE	0.0	0.0				

5.8.5 Defining Storage Pool Volumes

The following shows how to set up both tape storage pool volumes for an automated library and disk storage pool volumes.

5.8.5.1 Tape

Setting up storage pool volumes for an automated library is done in three stages:

1. Label volumes and check them into scratch pool
2. Check for outstanding mount requests
3. Reply to mount requests and issue labels

There are two methods of labelling volumes and checking them into the scratch pool remotely: using the Web administrative client, and using the administrative command line interface.

Both methods require checking for outstanding requests and issuing a reply to the request number with the label. If you have a library with a bar code reader, this simplifies the process greatly. Next you need to change the LABELSOURCE parameter in the `label libvolume` command from PROMPT to BARCODE.

Both methods require several stages to complete the process. You may find it easier to do this from the server itself, using the DSMLABEL utility and the `checkin libvolume` command. In this case, you would also be on hand to physically add and remove volumes from the library.

Label Volumes and Check Them In

The following command searches for tapes in the library named LDLT that are not defined to ADASM, and prompts you for a label. It does not overwrite any tapes that have a label already, and it does not relabel any tapes already checked in.

```

adsm> label libvolume ldlt search=yes labelsorce=prompt checkin=scratch \
cont> overwrite=no
Session established with server AIX1: AIX-RS/6000
Server Version 3, Release 1, Level 2.1
Server date/time: 02/04/99 16:46:07 Last access: 02/04/99 16:26:22

ANS8003I Process number 7 started.

```

Check for Outstanding Mount Requests

To see if there are any pending mount requests, use the `query request` command.

```

adsm> query request
ANR8352I Requests outstanding:
ANR8809I 006: Please provide the label name for the volume in slot element 8 of
library LDLT by issuing REPLY n LABEL=xxx within 51 minutes, where n is the
request ID and xxx is the desired label name.

```

Or, you can check the activity log to find out the number to reply to this request.

```

adsm> query actlog begintime=-00:01

```

Date/Time	Message
02/04/99 18:08:26	ANR2017I Administrator ADMIN issued command: QUERY ACTLOG -BEGINTime=-00:10
02/04/99 18:08:27	ANR8809I 006: Please provide the label name for the volume in slot element 8 of library LDLT by issuing REPLY n LABEL=xxx within 51 minutes, where n is the request ID and xxx is the desired label name.

Here we see the ANR8809I message, which is followed by the request number (006).

Note:

You also need to complete this step when using the Web administrative client.

Reply to Mount Requests and Issue Labels

To label a tape we need to reply to this message. The following command replies to request number 6 and issues a label name of V00001 to the tape volume.

```

adsm> REPLY 6 LABEL=v00001
ANR8499I Command accepted.

```

ADSM continues searching through the library for non-defined tapes and continues issuing requests for labels. For this example, we had two volumes, and cycling through the procedure above, we labelled and checked both volumes into the scratch pool.

```
adsm> query libvolume
```

Library Name	Volume Name	Status	Last Use	Home Element
LDLT	V00001	Scratch		8
LDLT	V00002	Scratch		9

Note:

If you label volumes incorrectly and wish to re label them, you will need to check the volume out of the library.

5.8.5.2 Disk

The size of the volumes you create depends on the values you decided are necessary in your initial planning. The values used below are small and are intended to be used only as an example.

When deciding how many volumes to use, consider this. It may be better to use a large volume to cover most of the expected storage pool size and then some smaller volumes to allow yourself flexibility in reassigning volumes to different storage pools as the need arises.

An example of this strategy could be to increase the size of the archive pool for monthly archives. In this case, you could migrate the backup random access storage pool to its next sequential storage pool. Redefine some of the backup storage pool volumes as archive pool volumes, complete the monthly archive, and back up the pool to offsite. You can then reassign volumes to their original state. This is just one example of how you can more efficiently use the resources at your disposal, since having a number of volumes in your storage pool allows you to redefine the volume or volumes to cover the extra capacity you need.

The process of adding a new storage pool volume consists of two steps:

1. Format the volume using the DSMFMT utility
2. Define the volume to ADSM using the `define volume` command

In a non-MVS environment, those two steps can be combined into a single process: you format the volume and immediately define it in ADSM by using the `define volume` command using the `FORMATSIZE` option. We recommend that you use the one-step method.

The following commands:

1. Format a data volume named `file01` under the `/adsm/stgpool` directory. The size is specified in MB and in this case is 200MB. It then assigns the volume to the disk storage pool `DISKDATA`.
2. Format a data volume named `dirs01` under the `/adsm/stgpool` directory of size 64MB. It then assigns the volume to the disk storage pool `DISKDIRS`.

```

adsm> define volume diskdata /adsm/stgpool/file01 formatsize=200
ANR2491I Volume Creation Process starting for /adsm/stgpool/file01, Process Id 39.

adsm> define volume diskdirs /adsm/stgpool/dirs01 formatsize=64
ANR2491I Volume Creation Process starting for /adsm/stgpool/dirs01, Process Id 40.

adsm> q vol

```

Volume Name	Storage Pool Name	Device Class Name	Estimated Capacity (MB)	Pct Util	Volume Status
/adsm/stgpool/dirs01	DISKDIRS	DISK	64.0	0.0	On-Line
/adsm/stgpool/file01	DISKDATA	DISK	200.0	0.0	On-Line

There is a guide for sizing the DISKDATA storage pool in the planning chapter, but the DISKDIRS storage pool must be treated differently. To size this pool depends largely on the types of directories you backup. NetWare directories, because of the many trustee directory assignments, take more space than Windows NT directories, which have fewer options for access permissions on their directory structure. At the other end of the scale, AIX directories without access control lists (ACLs) may all be saved as part of the ADSM database, and will not increase the space utilization percentage of the DISKDIRS storage pool.

We recommend that you initially choose a capacity of 64MB or 128MB. The OFFDIRS storage pool is also small and is likely to use part of one volume. However, the advantages of fast restores—in case of disaster—outweigh this small inefficiency.

5.8.6 Deleting Default Storage Pools

You can now delete the default storage pools to tidy up your installation, by using the `delete stgpool` command:

```

adsm> delete stgpool backuppool
ANR2201I Storage pool BACKUPOOL deleted.

adsm> delete stgpool archivepool
ANR2201I Storage pool ARCHIVEPOOL deleted.

adsm> delete stgpool spacemgpool
ANR2201I Storage pool SPACEMGPOOL deleted.

adsm> delete stgpool diskpool
ANR2201I Storage pool DISKPOOL deleted.

```

In 1.2, “Redbook Support Material” on page 2, we show how to load a predefined macro into ADSM. The macro `mac.stgdelete`, which we provide to delete the default storage pools, is shown in Appendix B.1.7, “Delete Default Storage Pools” on page 249. This macro deletes the default storage pools automatically set up by the Windows NT installation; these include a storage pool named DISKPOOL. (An ADSM install on a UNIX platform does not create this pool, so this line needs to be removed before running the macro.)

5.9 Additional Commands

In the following we provide some examples of commands that have a direct relationship with data storage.

5.9.1 Audit Library

The `audit library` command verifies that an automated library is in a consistent state with respect to the server's internal library volume inventory. This command does not start until all volumes in the library are dismounted. You can force this using the `dismount volume` command. It then reads the labels from all the volumes in the library. This can be the label written on the media itself (default), or you can indicate a barcode label by specifying `CHECKLABEL=BARCODE`.

The following command audits the library named LDLT:

```
adsm> audit library ldlt
ANS8003I Process number 40 started.

adsm> query process

Process Process Description  Status
Number
-----
      40 AUDIT LIBRARY          ANR8459I Auditing volume inventory for library
                                LDLT.

adsm> query actlog begintime=-00:05

Date/Time      Message
-----
02/11/99  04:21:12 ANR2017I Administrator ADMIN issued command: AUDIT LIBRARY
                                ldlt
02/11/99  04:21:12 ANR0984I Process 40 for AUDIT LIBRARY started in the
                                BACKGROUND at 04:21:12.
02/11/99  04:21:12 ANR8457I AUDIT LIBRARY: Operation for library LDLT started
                                as process 40.
02/11/99  04:24:43 ANR8461I AUDIT LIBRARY process for library LDLT completed
                                successfully.
02/11/99  04:24:43 ANR0985I Process 40 for AUDIT LIBRARY running in the
                                BACKGROUND completed with completion state SUCCESS at
                                04:24:43.
```

5.9.2 Audit Volume

The `audit volume` command checks for inconsistencies between the database references and the volume.

The following command audits a volume named V00003 and logs discrepancies to the activity log.

```
adsm> audit volume v00003 fix=no
ANR2310W This command will compare all inventory references to volume V00003
with the actual data stored on the volume and will report any discrepancies;
the data will be inaccessible to users until the operation completes.

Do you wish to proceed? (Yes/No) yes
ANR2313I Audit Volume (Inspect Only) process started for volume V00003 (process
ID 41).
ANS8003I Process number 41 started.
```

Using parameter FIX=NO reports any discrepancies to the activity log only. No changes are made. The course of action you take, if problems are found, may differ, so this monitoring and reporting process should be used first.

```

adsm> query process

Process Process Description  Status
Number
-----
41 Audit Volume             Volume V00003 (storage pool TAPEDATA), Files
  (Inspect Only)           Processed: 5001, Damaged Files Found: 0, Partial
                           Files Skipped: 0. Current Physical File (bytes):
                           1,981,318

                           Current input volume: V00003.

adsm> query actlog begintime=-00:08

Date/Time      Message
-----
02/11/99  04:30:51  ANR2017I Administrator ADMIN issued command: AUDIT VOLUME
v00003 Fix=NO
02/11/99  04:30:51  ANR1199I Removable volume V00003 is required for audit
process.
02/11/99  04:30:51  ANR8324I DLT volume V00003 is expected to be mounted
(R/W).
02/11/99  04:30:51  ANR0984I Process 41 for AUDIT VOLUME (INSPECT ONLY)
started in the BACKGROUND at 04:30:51.
02/11/99  04:30:51  ANR2313I Audit Volume (Inspect Only) process started for
volume V00003 (process ID 41).
02/11/99  04:32:11  ANR8337I DLT volume V00003 mounted in drive DRV1
(/dev/mt1).
02/11/99  04:37:11  ANR2315I Audit volume process ended for volume V00003;
9738 files inspected, 0 damaged files found and marked as
damaged.
02/11/99  04:37:11  ANR0987I Process 41 for AUDIT VOLUME (INSPECT ONLY)
running in the BACKGROUND processed 9738 items with a
completion state of SUCCESS at 04:37:11.

```

5.9.3 Back Up Storage Pool

The backup `stgpool` command backs up a primary storage pool to a copy pool. This command, to back up from one sequential storage pool to another, requires two physical drives. The two drives can be part of the same library. To back up a random access storage pool to a sequential storage pool requires only one drive.

The following command backs up a storage pool named TAPEDATA to a copy pool named OFFDATA.

```

adsm> backup stgpool tapedata offdata
ANS8003I Process number 45 started.

adsm> query process

Process Process Description  Status
Number
-----
45 Backup Storage Pool      Primary Pool TAPEDATA, Copy Pool OFFDATA, Files
Backed Up: 1, Bytes Backed Up: 25,049,183,
Unreadable Files: 0, Unreadable Bytes: 0.
Current Physical File (bytes): 26,114,266

Current input volume: V00003.

Current output volume: VP9999.

```

Check the activity log to see the successful completion, as shown in this example.

```

adsm> query actlog
-----
Date/Time      Message
-----
02/11/99  05:55:02  ANR1212I Backup process 45 ended for storage pool
              TAPEDATA.
02/11/99  05:55:02  ANR0986I Process 45 for BACKUP STORAGE POOL running in the
              BACKGROUND processed 9738 items for a total of
              649,109,355 bytes with a completion state of SUCCESS at
              05:55:02.
02/11/99  05:55:20  ANR1214I Backup of primary storage pool TAPEDATA to copy
              storage pool OFFDATA has ended. Files Backed Up: 9738,
              Bytes Backed Up: 649109355, Unreadable Files: 0,
              Unreadable Bytes: 0.

```

5.9.4 Checkin Libvolume

The `checkin libvolume` command checks in a library volume physically placed in the library, so that it can be seen by ADSM.

The following command searches the library named LDLT for new volumes. Once the volume is found, it is labelled V00001 and checked into the library as a scratch volume.

```

adsm> checkin libvolume ldlt v00001 search=yes status=scratch
ANS8003I Process number 39 started.

adsm> query libvolume
-----
Library Name  Volume Name  Status          Last Use      Home Element
-----
LDLT          V00001      Scratch         8
LDLT          V00003      Private         Data          5

```

5.9.5 Checkout Library Volumes

There are three parts to this operation:

- Check out the library volume.
- Check for outstanding mount requests.
- Reply to mount request.

Checkout Library Volumes

This command checks out library volume named V00002 from a library named LDLT.

```

adsm> checkout libvolume ldlt v00002
ANS8003I Process number 12 started.

adsm> query libvolume
-----
Library Name  Volume Name  Status          Last Use      Home Element
-----
LDLT          V00001      Scratch         8

```

Check for Mount Requests

To check for outstanding mount requests, issue a `query request` command or check the activity log.

```
adsm> query request
ANR8352I Requests outstanding:
ANR8307I 013: Remove DLT volume V00002 from slot with element number 9 of
library LDLT; issue 'REPLY' along with the request ID when ready.
```

Reply to Mount Request

Issue the `reply` command.

The `reply` command replies to request number 013 to confirm that volume V00002 has been removed from the library. To relabel the tape only, there is no need to physically remove the volume from the library.

```
adsm> reply 013
ANR8499I Command accepted.

adsm> query libvolume
```

Library Name	Volume Name	Status	Last Use	Home Element
LDLT	V00001	Scratch		8

Relabelling an Incorrectly Labelled Tape

If you wish to relabel an incorrectly labelled tape, ensure that parameter `OVERWRITE=YES` is set when executing the `label libvolume` command. The following command searches for and relabels any library volumes not yet checked in to AD SM, and overwrites an existing label.

```
adsm> label libvolume ldlt search=yes labelsource=prompt checkin=scratch \
cont> overwrite=yes
ANS8003I Process number 20 started.
```

If you want to relabel just one volume, you may specify it explicitly rather than use the `SEARCH=YES` and the `LABELSOURCE=PROMPT` parameter, but you still need to check it out first.

5.9.6 Delete a Library

To explain the delete command, we have an example for deleting a library. This shows how the different levels are interlinked and must be considered for successful deletion.

To delete a library you have to remove any references to it. There are many levels above a library that you may want deleted or updated to point to something else. You need to delete the following objects in the correct order:

1. Storage pool
2. Device class
3. Drive
4. Library

The following example explains how to delete library M8MM we set up in 5.7.1, “Defining a Library — Physical” on page 77. We show how to delete the whole structure, including the storage pool initially set up for this library.

5.9.6.1 Delete Storage Pool

Before you delete the storage pool TAPEDATA, you need to remove any references to it. Check the hierarchical structure using `query stgpool` commands.

```
adsm> query stgpool diskdata
```

Storage Pool Name	Device Class Name	Estimated Capacity (MB)	Pct Util	Pct Migr	High Mig Pct	Low Mig Pct	Next Storage Pool
DISKDATA	DISK	100.0	0.0	0.0	70	30	TAPEDATA

```
adsm> query stgpool tapedata
```

Storage Pool Name	Device Class Name	Estimated Capacity (MB)	Pct Util	Pct Migr	High Mig Pct	Low Mig Pct	Next Storage Pool
TAPEDATA	C8MM	0.0	0.0	0.0	100	99	

Update storage pool DISKDATA to remove the reference to TAPEDATA as its next storage pool.

```
adsm> update stgpool diskdata nextstgpool=""
ANR2202I Storage pool DISKDATA updated.
```

```
adsm> query stgpool diskdata
```

Storage Pool Name	Device Class Name	Estimated Capacity (MB)	Pct Util	Pct Migr	High Mig Pct	Low Mig Pct	Next Storage Pool
DISKDATA	DISK	100.0	0.0	0.0	70	30	

The following command deletes the storage pool named TAPEDATA.

```
adsm> delete stgpool tapedata
ANR2201I Storage pool TAPEDATA deleted.
```

Note:

You can only delete a storage pool if there are no storage pool volumes assigned to it.

5.9.6.2 Delete Device Class

Next you have to delete the device class of the library. The following commands query all defined device classes and then delete the device class named C8MM.

```
adsm> query devclass
```

Device Class Name	Device Access Strategy	Storage Pool Count	Device Type	Format	Est/Max Capacity (MB)	Mount Limit
C8MM	Sequential	0	8MM	DRIVE	0.0	DRIVES
DISK	Random	4				

```
adsm> delete devclass c8mm
ANR2204I Device class C8MM deleted.
```

5.9.6.3 Delete Drive

Next you have to delete the drive of the library. The following commands query all defined drives and delete the drive named DRV5 from library named M8MM.

```
adsm> query drive
```

Library Name	Drive Name	Device Type	Device	ON LINE
M8MM	DRV5	8MM	MT5.0.0.2	Yes

```
adsm> delete drive m8mm drv5
ANR8412I Drive DRV5 deleted from library M8MM.
```

5.9.6.4 Delete Library

Now library M8MM can be deleted.

```
adsm> delete library m8mm
ANR8410I Library m8mm deleted.

adsm> query library
ANR2034E QUERY LIBRARY: No match found using this criteria.
ANS8001I Return code 11.
```

5.9.7 Dismount a Volume

The following command dismounts the volume named V00003 that is physically inside one of the library's drives.

```
adsm> dismount volume v00003
ANR8499I Command accepted.

adsm> query mount
ANR8331I DLT volume V00003 is mounted R/W in drive DRV1 (/dev/mt1), status:
DISMOUNTING.
ANR8334I          1 volumes found.

adsm> query mount
ANR2034E QUERY MOUNT: No match found using this criteria.
```

5.9.8 Migration

Migration is forced by updating the threshold parameters for migration on the storage pool to a value that causes migration to start. To clear all data from one pool to the next storage pool, you change the high and low values to 0.

Check storage pool migration values before you change them.

```
adsm> query stgpool
```

Storage Pool Name	Device Class Name	Estimated Capacity (MB)	Pct Util	Pct Migr	High Mig Pct	Low Mig Pct	Next Storage Pool
DISKDATA	DISK	50.0	43.1	43.1	70	30	TAPEDATA
TAPEDATA	CDLT	40,960.0	1.5	50.0	100	70	

The following command forces migration for storage pool DISKDATA. This is done by setting both storage pool parameters, the high threshold and the low threshold, to 0. Use the query process command to check on the progress. You also can search the activity log over the last hour for any string "migra".

```
adsm> update stgpool diskdata high=0 low=0
ANR2202I Storage pool DISKDATA updated.
```

```
adsm> query process
```

Process Number	Process Description	Status
51	Migration	Disk Storage Pool DISKDATA, Moved Files: 62, Moved Bytes: 11,603,968, Unreadable Files: 0, Unreadable Bytes: 0. Current Physical File (bytes): 2,068,480 Current output volume: V00003.

```
adsm> query actlog search=migra*
```

Date/Time	Message
02/11/99 23:46:48	ANR0984I Process 51 for MIGRATION started in the BACKGROUND at 23:46:48.
02/11/99 23:46:48	ANR1000I Migration process 51 started for storage pool DISKDATA.
02/11/99 23:48:58	ANR1001I Migration process 51 ended for storage pool DISKDATA.
02/11/99 23:48:58	ANR0986I Process 51 for MIGRATION running in the BACKGROUND processed 145 items for a total of 22,589,440 bytes with a completion state of SUCCESS at 23:48:58.

The data has moved from the storage pool DISKDATA, to storage pool TAPEDATA. After migration has completed, you have to change the threshold values back to their original levels.

```

adsm> query stgpool

Storage Pool Name   Device Class Name   Estimated Capacity (MB)   Pct Util   Pct Migr   High Mig Pct   Low Mig Pct   Next Storage Pool
-----
DISKDATA            DISK                50.0                    0.0        0.0        0          0          0          0          TAPEDATA
TAPEDATA            CDLT                40,960.0                1.6        50.0       100        70

```

```

adsm> update stgpool diskdata high=70 low=30
ANR2202I Storage pool DISKDATA updated.

```

5.9.9 Move Data

The `move data` command moves the data from one volume to another within the same storage pool (this requires the storage pool to have two drives), or to another storage pool.

The following command moves the data from volume V00003 in the TAPEDATA storage pool to a volume in the 8MMPOOL.

```

adsm> move data v00003 stgpool=8mmpool
ANR2233W This command will move all of the data stored on volume V00003 to other volumes in storage pool 8MMPOOL; the data will be inaccessible to users until the operation completes.

Do you wish to proceed? (Yes/No) yes
ANS8003I Process number 48 started.

```

Here is an example where you may use this command: A volume showed a write error, but after checking for any inconsistencies between the volume and the database, the data on the volume was determined to be good. However, you do not want to use this volume any more. You may now move all the data from this volume to another volume in the same storage pool. The volume with an error can now be checked out of the library and discarded.

5.9.10 Query Content

The `query content` command queries the contents of a volume and lists all files contained therein. The output for this command can be very long.

```

adsm> query content v00003

Node Name           Type Filespace   Client's Name for File
-----
ANDY                Bkup \\23831vn- \ADDITIONAL\ADSM\CLIENT\ DATA1.CAB
                   \c$
ANDY                Bkup \\23831vn- \ADDITIONAL\ADSM\CLIENT\ IP21494.EXE
                   \c$
ANDY                Bkup \\23831vn- \ADDITIONAL\ADSM\CLIENT\
                   \c$ IP21494_UNPACK.BAT

```

It can be reduced by specifying the number of files you want displayed by using the `COUNT` parameter. The following command lists the contents of volume V00003, but restricts the output to files that are backed up to a copy storage pool, and only shows four lines of output. If the parameter was `COPIED=NO`, the output would only show files that were not backed up to a copy pool.

```
adsm> query content v00003 copied=yes count=4
```

Node Name	Type	Filespace Name	Client's Name for File
ANDY	Bkup	\\23831vn-\c\$	\ADDITIONAL\ADSM\CLIENT\ DATA1.CAB
ANDY	Bkup	\\23831vn-\c\$	\ADDITIONAL\ADSM\CLIENT\ IP21494.EXE
ANDY	Bkup	\\23831vn-\c\$	\ADDITIONAL\ADSM\CLIENT\ IP21494_UNPACK.BAT
ANDY	Bkup	\\23831vn-\c\$	\ADDITIONAL\ADSM\CLIENT\ LANG.DAT

5.9.11 Query Occupancy

The following command displays information on where a client's filespace are stored and how much space they occupy.

```
adsm> query occupancy andy
```

Node Name	Type	Filespace Name	Storage Pool Name	Number of Files	Physical Space Occupied (MB)	Logical Space Occupied (MB)
ANDY	Bkup	\\23831vn-\c\$	BACKUPPOOL	639	1.77	1.77
ANDY	Bkup	\\23831vn-\c\$	OFFDATA	9,738	619.04	619.04
ANDY	Bkup	\\23831vn-\c\$	TAPEDATA	9,883	640.54	640.54

The query occupancy command can also be used to determine whose filespace are occupying a certain device class or storage pool.

```
adsm> query occupancy devclass=disk
```

Node Name	Type	Filespace Name	Storage Pool Name	Number of Files	Physical Space Occupied (MB)	Logical Space Occupied (MB)
AIXCLIENT1	Bkup	/home	BACKUPPOOL	13	0.17	0.17
AIXCLIENT1	Arch	/home	BACKUPPOOL	13	0.17	0.17
ANDY	Bkup	\\23831vn-\c\$	BACKUPPOOL	639	1.77	1.77
SCANDIUM	Arch	CDRIVE	BACKUPPOOL	163	543.69	543.69

5.9.12 Rename a Storage Pool

The rename stgpool command renames a storage pool. The following command renames SPACEMGPOOL to FUTUREHSM.

```
adsm> rename stgpool spacemgpool futurehsm
ANR2213I RENAME STGPOL: Storage pool SPACEMGPOOL renamed to FUTUREHSM.
```

5.9.13 SQL Commands

ADSM uses a database that accepts structured query language (SQL) commands. ADSM provides three system catalog tables:

- SYSCAT.TABLES: Contains information about all tables that can be queried.
- SYSCAT.COLUMNS: Describes the columns in each table.
- SYSCAT.ENUMTYPES: Defines the valid values for each enumerated data type.

There are many commands that can be generated in this way to display a customized query from the ADSM database. There are a small number of useful commands listed here to start with.

Example 1

The following command lists all volumes from table volumes where they are in an error condition. In our example, this command fails, since we do not have any volumes in an error condition.

```
adsm> select * from volumes where error_state='YES'  
ANR2034E SELECT: No match found using this criteria.  
ANS8001I Return code 11.
```

Example 2

The following command counts the number of volumes within each storage pool. This can be useful where large libraries with many volumes are concerned.

```
adsm> select count(*),stgpool_name from volumes group by stgpool_name  
  
Unnamed[1]      STGPOOL_NAME  
-----  
1              8MMPPOOL  
2              BACKUPPOOL  
1              DISKDATA  
1              DISKDIRS  
1              OFFDATA  
1              TAPEDATA
```

Example 3

The following command lists how many files have been backed up from each node. If you replace BACKUPS with ARCHIVES you see the number of files archived from each node.

```
adsm> select node_name,count(*) from backups group by node_name  
ANR2963W This SQL query may produce a very large result table, or may require a  
significant amount of time to compute.  
  
Do you wish to proceed? (Yes/No) yes  
  
NODE_NAME      Unnamed[2]  
-----  
AIXCLIENT1    57  
ANDY          10522  
NOGALLES_EXCH 2  
PAGOPAGO_EXCH 13  
SCANDIUM_EXCH 1  
WKS.23FFHBV  51
```

Example 4

The following command shows the volumes in which all filespace from client AIXCLIENT1 are residing on V00002.

```
adsm> select distinct node_name,volume_name from volumeusage where \  
cont> node_name='AIXCLIENT1' and stgpool_name='TAPEDATA'
```

NODE_NAME	VOLUME_NAME
AIXCLIENT1	V00002

For further information on SQL commands, use the `help select` command.

There are also some useful sample scripts available. See 3.4.1, “Load Sample Scripts” on page 42 for an explanation of how to load the sample scripts into the ADSM database.

Chapter 6. Data Storage Policy

A data storage environment consists of three types of resources: machines, rules, and data. The machines are computers containing data that must be backed up and the rules specify how the backup copies are to be treated. Basically, a data storage policy defines the relationships between these three resources. Figure 12 on page 105 illustrates this policy relationship. See 2.2, “Data Retention Requirements” on page 15 for further planning considerations.

Depending on how you need to manage your backup data, your data storage policy can be very simple or very complex. The simplest policy would have one set of rules that apply to all of the data in your organization while the most complex policy would be to have a unique set of rules for each individual file. Most installations fall somewhere between these two extremes.

ADSM has entities that group and organize the resources and define relationships between them. A machine, or node in ADSM terminology, is grouped together with other nodes into a policy domain. The domain links the nodes to a policy set which consists of management classes. A management class contains rules called copy groups that it links to the data. When the data is linked to particular rules, it is said to be bound to the management class that contains the rules.

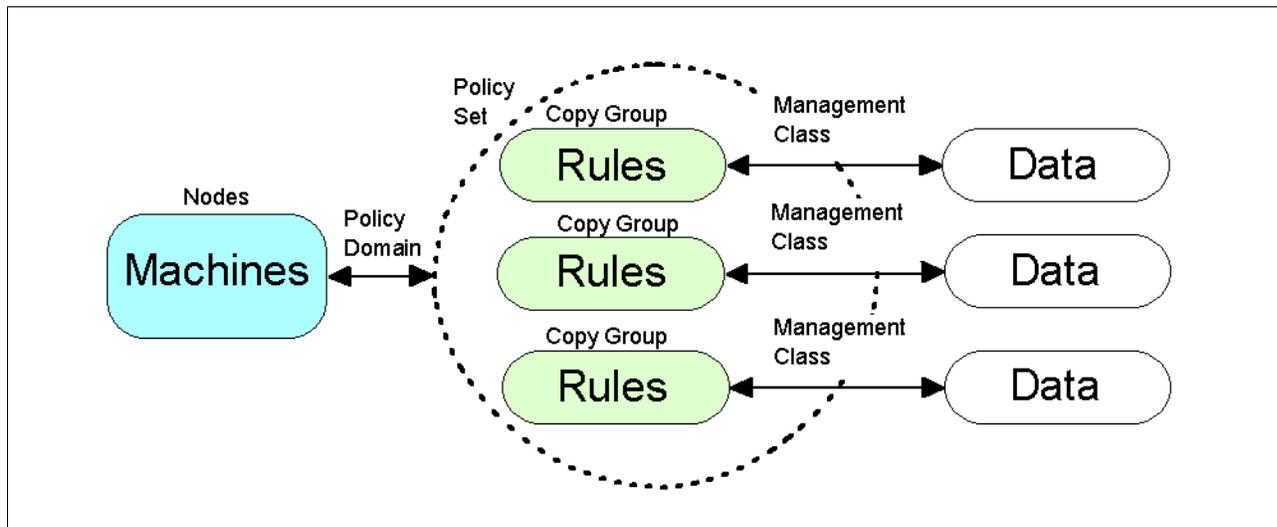


Figure 12. Policy Relationships and Resources

Another way to look at the components that make up a policy is to consider them in the hierarchical fashion in which they are defined. That is, consider the policy domain containing the policy set, the policy set containing the management classes, and the management classes containing the copy groups. Our view of the policy components is illustrated in Figure 13 on page 106.

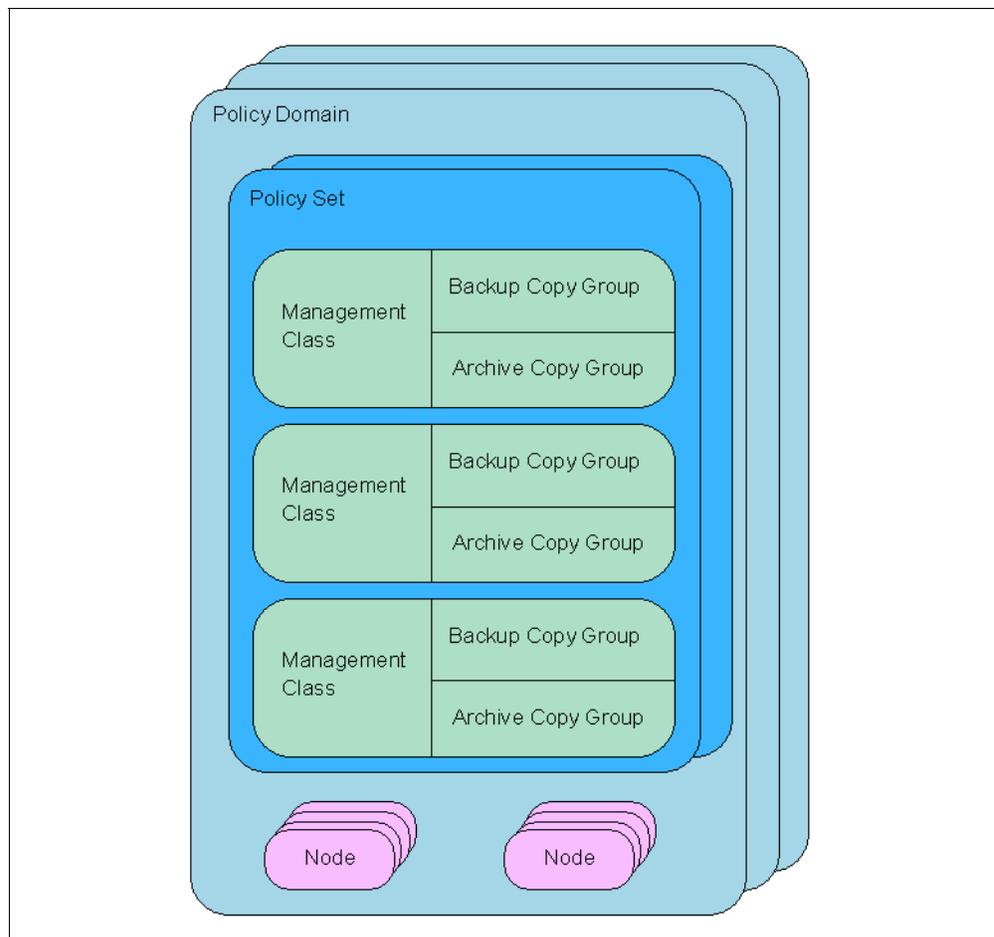


Figure 13. Data Storage Policy Components

6.1 Setting Up Data Storage Policies

This chapter takes you through the steps necessary to define a policy. The minimum steps required are:

- Define a policy domain
- Define a policy set
- Define a management class
- Define a backup copy group
- Define an archive copy group

ADSM has a default policy, called STANDARD, that is configured when the product is installed. The STANDARD policy is good for product evaluation and testing, but the recommended configuration we have developed for this redbook is a better starting point. It is fully integrated with the rest of the ADSM server components to get your system up and running as quickly as possible.

6.1.1 Defining a Policy Domain

A policy domain is a way to group ADSM clients depending on how you want to treat their data. It allows you to logically group the machines in your organization according to:

- **Default policy:** The default set of rules to apply to the clients. The rules define the storage management policy, including how many copies of data to keep and how long to keep them.
- **Administrative control:** Access to the clients and their policy can be restricted to certain administrators.

Let us consider a typical organization consisting of UNIX and Windows NT machines. The UNIX machines are large file servers and need many copies of their data maintained for a long period of time. The UNIX support group is the only group authorized to access the UNIX machines. The UNIX policy domain would hold all of the UNIX machines, and would only be accessible to the UNIX ADSM administrators. The default policy in this domain would apply only to the UNIX machines.

The Windows NT machines are workstations and need a few copies of their data maintained for a short period of time. The Windows NT support group is the only group qualified to access the Windows NT machines. The Windows NT policy domain would hold all of the Windows NT machines, and access to it would be restricted to the Windows NT support group. The default policy in this domain would only apply to the Windows NT machines.

This is a good example of how default policy and administrative control can be used to break up your organization into policy domains.

The `define domain` command creates new policy domains. The following command shows how to create a policy domain called `production` which is used to group the production servers for the ABC corporation.

```
adsm> define domain production backret=100 archret=365 \  
description="Production servers for ABC Corporation"  
ANR1500I Policy domain PRODUCTION defined.
```

The ADSM clients that belong to this domain share the same default policy, and access to them can be restricted to administrators authorized to manage the production servers. The nodes must be assigned to this policy domain when they are registered with the ADSM server as described in Chapter 7.2.3.1, “Registering a Client Node” on page 128.

The `define domain` command also defines two parameters that are used in special cases to determine how data is handled:

- **BACKRET:** How long backup data is retained if it is not governed by the existing backup policies
- **ARCHRET:** How long archive data is retained if it is not governed by the existing archive policies

These two retention grace period parameters are used as a safety net for data that is ‘orphaned’ from a valid set of policies for some reason. This can happen if

a set of policies were accidentally deleted while still containing data. The BACKRET and ARCHRET parameters are used to manage this data.

6.1.2 Defining a Policy Set

The policy set is a grouping of management classes. There can be multiple policy sets within a policy domain, but only one policy set can be active at a time. This restriction has resulted in most installations using only one policy set per policy domain.

The ACTIVE policy set is a special entity in the policy domain. It exists in every domain and cannot be changed directly. To change the ACTIVE policy set, you must define your rules in a policy set which is subsequently validated and activated. The activation process takes a snapshot of the policy set and places it in the ACTIVE policy set. It is important to note that the ACTIVE policy set is a point-in-time snapshot of the originating policy set. Therefore, further changes to the originating policy set have no effect on the ACTIVE policy set until the changed policy set is validated and activated.

The `define policyset` command creates a new policy set within a policy domain. Although the policy set can be named anything except the reserved name ACTIVE, it is common practice to give the same name to the policy domain and policy set for simplicity. The following command shows how to define a PRODUCTION policy set for the PRODUCTION policy domain.

```
adsm> define policyset production production \  
description="Production servers for ABC Corporation"  
ANR1510I Policy set PRODUCTION defined in policy domain PRODUCTION.
```

6.1.3 Defining Management Classes

A management class contains specific rules and is the link between rules and data. The rules are contained in constructs called copy groups, one for rules governing backup data and one for rules governing archive data. A management class can have a backup copy group or an archive copy group or both.

A policy set may contain many management classes but only one is designated as the default. The default management class in the policy set is linked to any data in the domain that is not explicitly linked to another management class.

Explicitly linking, or binding, files to management classes can be accomplished in three different ways depending on the type of data. Backup data is bound to a management class using the INCLUDE option of the ADSM client's include-exclude list. Archive data is bound to a management class using the ARCHMC command line option. Directory data can be bound to a management class using the DIRMC client option.

Figure 14 on page 109 shows an example of an INCLUDE statement that assigns the backup file `/home/admin/redbook.script` to a management class called REDBOOK while allowing the rest of the backup files in `/home/admin` to go to the default management class. The binding is actually done during the backup operation. See , "Include-Exclude Options" on page 155 and Appendix B.3, "Client Options Files" on page 254 for more examples of the INCLUDE statement.

```
include /home/admin/.../*
include /home/admin/redbook.script redbook
```

Figure 14. Include Option Example

Figure 15 on page 109 shows an example of the ARCHMC command line option. The file `/home/admin/redbook.doc` will be bound to the management class REDBOOKARCHIVE. Without the ARCHMC option, this data would have been bound to the default management class. See Chapter 10.3.1.4, “Running Archive Operations” on page 163 for more examples of the ARCHMC option.

```
dsmc archive -archmc=redbookarchive /home/admin/redbook.doc
```

Figure 15. ARCHMC Example

Figure 16 on page 109 shows an example of the DIRMC option in a client options file. All of the directories from this client will be bound to the directory management class. Make sure that the retention period for directories is set longer than for the data itself. Without the DIRMC option, the directory information would go to the management class with the longest retention period. You can use this option to make your restores quicker by keeping your directory data on a primary disk storage pool instead of on tape. The directory information is the first thing that is rebuilt during a restore and by making it quickly accessible, your restore time will be improved. The recommended policy configuration includes a directory management class for the directory information, and the DIRMC option is specified in the recommended client option sets described in Appendix B.3, “Client Options Files” on page 254.

```
dirmc directory
```

Figure 16. DIRMC Example

The `define mgmtclass` command creates a new management class in the policy set and domain. The `assign defmgmtclass` command defines a management class as the default for the domain. The following commands show how to create the CUSTDATA management class in the PRODUCTION policy set within the PRODUCTION policy domain and how to assign it as the default management class.

```
adsm> define mgmtclass production production custdata \  
description="Custdata management class for ABC Corporation"  
ANR1520I Management class CUSTDATA defined in policy domain PRODUCTION, set  
PRODUCTION.  
  
adsm> assign defmgmtclass production production custdata  
ANR1538I Default management class set to CUSTDATA for policy domain PRODUCTION,  
set PRODUCTION.
```

6.1.4 Defining Backup Copy Groups

The backup copy group is concerned with two logical objects: the file, and the file copy. A file is the actual data on a node, while a file copy is a point-in-time copy of the file. Another way to think of it is that ADSM contains file copies, and nodes contain files.

A file can be in one of two possible states: existing or deleted. When we talk about an existing file on a node, we mean a file that has been previously backed up and still exists on the node. A deleted file is a file that has been previously backed up and has been deleted from the node. This simple concept is important when discussing data storage rules.

A file copy can be in one of three states: active, inactive, or expired. An active file copy is the most current copy of the file, an inactive file copy is a previous copy of the file, and an expired file copy is a copy to be removed from the ADSM server. A backup file copy is set to the expired state when it no longer conforms to the rules set forth in the backup copy group.

Whether the file exists or is deleted, the file copy always passes through the same states in the same order. A file copy starts out as active, since it is the first copy of the file, and therefore the most current. Once the file changes and we take another file copy, the first file copy changes to be inactive because we have a more recent one. Eventually, the first file copy expires based on one of two limits placed on it by our rules: number of copies or retention period.

The number of copies that we set in our rules specifies the total number of file copies to maintain in the ADSM database. It is important to note that the specified number includes the active file copy. Thus, when we set the number of file copies to three, we are keeping one active copy and two inactive copies. When the number of copies is exceeded, the oldest copies are removed from the database.

The retention period that we set in our rules specifies the length of time that we retain inactive file copies. Note that there is no retention period for active file copies; they exist as long as the file exists on the node.

Whether or not the file exists on the node affects which rules are used to expire the file copies. If the file exists, the following two backup copy group parameters are in effect:

- **VEREXISTS:** Specifies the number of file copies, or versions, to keep. This number includes active and inactive file copies.
- **RETEXTRA:** Specifies how long to keep inactive file copies. When a file changes from active to inactive, it is kept for these extra days and then removed. It is important to note that the retention period starts from when the file copy changes to inactive, and not from its original backup date.

If the file has been deleted, the active file copy is made inactive. At this point, there are only inactive file copies for this data in the ADSM server, and the following parameters apply:

- **VERDELETED:** Specifies the number of file copies to keep after the file has been deleted.
- **RETEXTRA:** Specifies how long to keep inactive file copies.

- **REONLY:** Specifies how long to maintain the last file copy of the data. This is the number of days to keep the last copy only, and does not apply to other inactive file copies which are still governed by the RETEXTRA parameter.

The backup copy group defines five other attributes that control the way that backup data is handled:

- **TYPE:** The TYPE parameter is used to differentiate between the two possible types of copy groups. In the case of a backup copy group, it is set to BACKUP.
- **DESTINATION:** The backup copy group specifies where to store the data sent to it from backup operations using the DESTINATION parameter. The copy group bridges the gap between data files and storage pools as illustrated in Figure 17 on page 111. The figure shows different types of data flowing through the copy groups and into the storage pools. Note that there is not necessarily a one-to-one relationship between copy groups and storage pools.

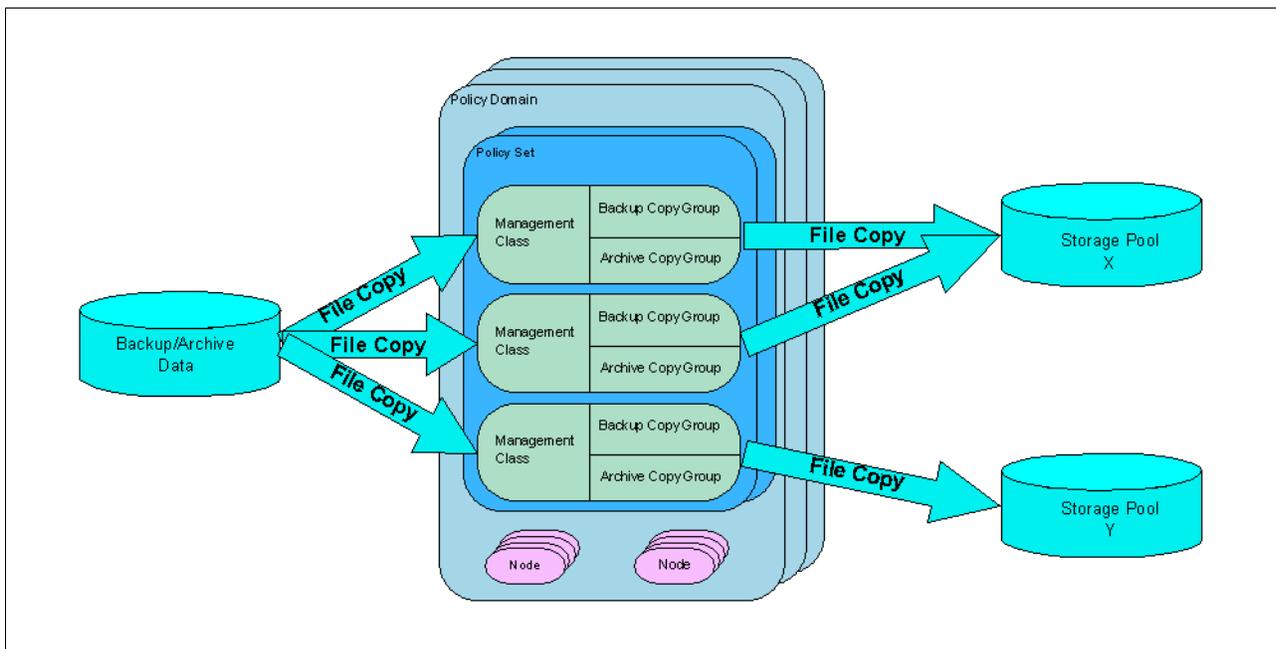


Figure 17. Data Flow through Copy Groups

- **MODE:** The MODE parameter specifies how files are to be selected for incremental backup. Setting the mode to MODIFIED allows a file to be backed up only if it has changed since the last backup. The ABSOLUTE setting allows files to be backed up, regardless of whether they have changed or not. The latter value would only be used for special cases; the default value is MODIFIED.
- **FREQUENCY:** The FREQUENCY parameter specifies how often to allow a file to be incrementally backed up. A selective backup, which backs up data regardless of whether it has changed or not, is not affected by this parameter. To incrementally back up a file from a node, three conditions must be satisfied:

1. Include-exclude statements allow the file to be considered for backup.
2. The file satisfies the MODE setting. That is, if the MODE is set to MODIFIED, the file must have changed to qualify for backup. If the MODE is set to ABSOLUTE, then the file is automatically allowed to be backed up.
3. The difference between the server time and the active file copy timestamp must be greater than FREQUENCY setting. The frequency is converted to hours to compare to the timestamp difference.

For example, take a file called `/home/admin/redbook.doc` that is eligible for backup in the include-exclude list, and that has changed since the last backup at 8 AM this morning. The server time is 11 AM when an incremental backup is started, so the difference between the server time and the file copy time is 3 hours. If the frequency is set to 1 day, then 24 hours must pass between incremental backups before a file is backed up again. Therefore, the `/home/admin/redbook.doc` file is not backed up, since 3 hours is less than 24 hours.

- **SERIALIZATION:** The SERIALIZATION parameter specifies what to do with files that are modified during a backup operation. When we say that a file is modified during backup, we mean that it is modified after ADSM examined it for its details, but before it was completely backed up to the server. This sort of backup is referred to as a "dirty backup" because the file is in an inconsistent state and may not restore properly. The SERIALIZATION parameter provides four options to deal with this problem:
 1. The SHRSTATIC setting specifies that a file is not backed up if it is modified during backup, but multiple attempts are made to back up the file. If the file is being modified through all of these attempts, the file is not backed up. The number of attempts can be controlled using the CHANGINGRETRIES option in the client options file.
 2. The STATIC setting specifies that a file is not backed up if it is modified during backup and no additional attempts are made.
 3. The SHRDYNAMIC setting specifies that a file is backed up if it is modified during backup, but multiple attempts are made to back it up without modification first. If that cannot be done, then the file is backed up anyway.
 4. The DYNAMIC setting specifies that a file is backed up even if it is modified during backup. There are no preliminary attempts to back up the file unmodified; it is backed up on the first attempt.

The `define copygroup` command creates a new copy group within the management class. Note that the command does not require a name for the backup copy group; it is always called STANDARD and cannot be changed. There is no reason to change this parameter, since it only has relevance within the ADSM server itself.

The following command shows how to create a backup copy group for the management class CUSTDATA in the policy set PRODUCTION of the policy domain PRODUCTION. The client checks for files that have been modified since the last backup process and tries to back them up to the DISKDATA primary storage pool. After a few retries, the process gives up on files that are being modified during backup. The client data is governed by the following rules:

- If the file exists on the client, then 3 versions of it are kept for 100 days.
- If the file no longer exists on the client, then the last copy is kept for 365 days.

```
adsm> define copygroup production production custdata type=backup \  
cont> destination=DISKDATA frequency=1 verexists=3 verdeleted=1 retextra=100 \  
cont> retonly=365 mode=modified serialization=shrstatic  
ANR1530I Backup copy group STANDARD defined in policy domain PRODUCTION, set  
PRODUCTION, management class CUSTDATA.
```

6.1.5 Defining an Archive Copy Group

An archive copy group consists of fewer parameters than the backup copy group. The most important of these are:

- **TYPE:** Parameter used to differentiate between two possible types of copy groups. In the case of an archive copy group, it is set to ARCHIVE.
- **DESTINATION:** Destination storage pool to use for data storage. Usually, this is the primary storage pool in a storage pool hierarchy including disk pools and tape pools.
- **RETVR:** How long to keep this archive copy. Usually it is set to 365 days.
- **SERIALIZATION:** How to deal with files that are being modified during processing. There are a few options for this parameter, but the only two that are of interest here are the SHRSTATIC and SHRDYNAMIC options. SHRSTATIC means that the process tries multiple times to backup a file that is in use and then gives up. SHRDYNAMIC means that the file is backed up even though it is still in use. This method is not the recommended way, but it should be used for files that are always in use and need to be archived.

Archive data is assigned to the archive copy group of the default management class unless you explicitly specify otherwise on the command line.

The `define copygroup` command is used to define the archive copy group. Note that the command does not require a name for the archive copy group; it is always called STANDARD and cannot be changed. There is no reason to change this parameter, since it only has relevance within the ADSTM server itself.

The following command shows how to define an archive copy group for the CUSTDATA management class in the PRODUCTION policy set within the PRODUCTION policy domain. The archive files are sent to the DISKDATA primary storage pool, and are kept for 365 days. If the operation cannot archive a file because it is in use, it tries a few times and then gives up.

```
adsm> define copygroup production production custdata type=archive \  
cont> destination=DISKDATA retver=365 serialization=shrstatic  
ANR1535I Archive copy group STANDARD defined in policy domain PRODUCTION, set  
PRODUCTION, management class CUSTDATA.
```

6.2 Recommended Setup

Figure 18 on page 114 is a graphical representation of the key components of our recommended policy configuration. We define two policy domains, SERVER and WORKSTN. Both domains have *similar* policy sets and management classes, but their copy group details show that the SERVER domain has more copies, and longer retention periods than the WORKSTN domain. It is not necessary to define multiple domains, but this demonstrates that the policy is really defined in the copy groups, and that the rest of the constructs are used primarily for flexibility.

Warning:

Care must be taken when migrating to another policy domain. If a client had data backed up under the old domain, the data is rebound to the same management class of the new domain, if it exists. If the management class does not exist in the new domain, the data is rebound to the default management class of the new domain. If this management class does not have a backup copy group, the data is rebound to the retention grace period defined in the new domain. This can lead to undesired expiration of backup versions.

In 1.2, “Redbook Support Material” on page 2, we show how to load a predefined macro into ADSM. The macro `mac.policy`, which we provide to define the recommended policy settings in our redbook environment, is shown in Appendix B.1.3, “Define Policy Structure” on page 245.

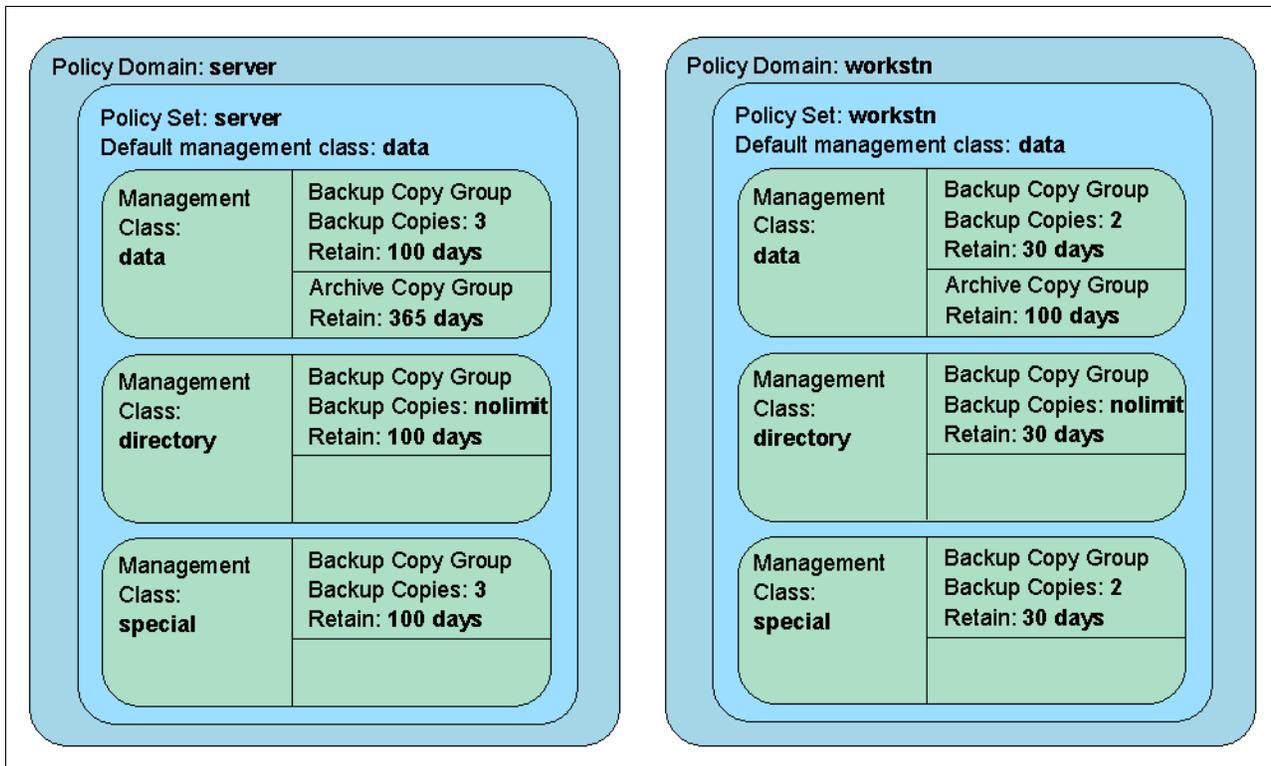


Figure 18. Recommended Policy Definitions

6.2.1 Defining Policy Domains

The recommended configuration consists of two policy domains that are used to separate server machines from workstation machines. The nature of the data found on each type of machine makes this a logical division of resources and follows the administrative boundaries used by most customers. This example defines the recommended policy domains:

```
adsm> define domain server description="Server nodes" \  
cont> backretention=100 archretention=365  
ANR1500I Policy domain server defined.  
  
adsm> define domain workstn description="Workstation nodes" \  
cont> backretention=100 archretention=365  
ANR1500I Policy domain workstn defined.
```

6.2.2 Defining Policy Sets

We recommend defining one policy set for each of the policy domains; for example:

```
adsm> define policyset server server description="Server nodes"  
ANR1510I Policy set server defined in policy domain server.  
  
adsm> define policyset workstn workstn description="Workstation nodes"  
ANR1510I Policy set workstn defined in policy domain workstn.
```

6.2.3 Defining Management Classes

We have defined three management classes for each domain in the recommended configuration. These management classes are used to categorize the domain data:

- **DATA:** This is the default management class for both domains. It is used for most of the data in the environment.
- **DIRECTORY:** Directory structure information is assigned to this class to make it possible to improve restore time.
- **SPECIAL:** This management class is used to store files that are modified during backup. An application log file could be bound to this management class to make sure that a copy is taken. A typical example of this type of file is an application log file that receives messages or errors. While it is being backed up, it is not going to be available for backup or archive unless the application is stopped to release the file. The management class specifications (in the backup/archive copy group definitions) include parameters for treatment of data that is being modified while a backup/archive operation is being performed. One of these options allows the file to be backed up even if it is being modified.

Space management clients are not covered in this redbook, but there is a space management parameter in the management class definition that we must change to avoid problems. The `MIGDESTINATION` parameter specifies the name of the storage pool for the migration of space managed files, and it defaults to a storage pool that we delete during the recommended configuration steps. We point it to a

dummy storage pool called NONE, which allows us to validate and activate the policy set without errors messages.

The management classes for the recommended configuration can be defined as follows:

```
adsm> define mgmtclass server server data migdestination=NONE \  
cont> description="Default management class for server domain"  
ANR1520I Management class data defined in policy domain server, set  
server.  
  
adsm> assign defmgmtclass server server data  
ANR1538I Default management class set to data for policy domain server, set  
server.  
  
adsm> define mgmtclass server server directory migdestination=NONE \  
cont> description="Directory management class for server domain"  
ANR1520I Management class directory defined in policy domain server, set  
server.  
  
adsm> define mgmtclass server server special migdestination=NONE \  
cont> description="Special management class for server domain"  
ANR1520I Management class special defined in policy domain server, set  
server.  
  
adsm> define mgmtclass workstn workstn data migdestination=NONE \  
cont> description="Default management class for workstn domain"  
ANR1520I Management class data defined in policy domain workstn, set  
workstn.  
  
adsm> assign defmgmtclass workstn workstn data  
ANR1538I Default management class set to data for policy domain workstn,  
set workstn.  
  
adsm> define mgmtclass workstn workstn directory migdestination=NONE \  
cont> description="Directory management class for workstn domain"  
ANR1520I Management class directory defined in policy domain workstn, set  
workstn.  
  
adsm> define mgmtclass workstn workstn special migdestination=NONE \  
cont> description="Special management class for workstn domain"  
ANR1520I Management class special defined in policy domain workstn, set  
workstn.
```

6.2.4 Defining Backup Copy Groups

The reasons for creating two domains (WORKSTN and SERVER) are clearly illustrated in the backup copy group parameters. The two domains differ in the following ways:

- VEREXISTS: The WORKSTN domain only maintains two copies of existing data, as compared to three copies in the SERVER domain.
- RETEXTRA: The WORKSTN domain only keeps extra copies for 30 days, as compared to 100 days in the SERVER domain.
- RETONLY: The WORKSTN domain keeps the last copy of deleted data for 100 days, as compared to 365 days in the SERVER domain.

An interesting parameter also illustrates the difference between the management classes:

- SERIALIZATION: The backup copy group within the special management class is set up to use SHR DYNAMIC instead of SHR STATIC. This management class should be assigned to log files and other cases where the files are always open.

Note:

If some of your data require policy settings different than our recommended defaults—for example, if you need to keep 7 copy versions of certain files—then you need to tailor an additional management class for your purpose, and assign the data to it using the INCLUDE option. See 10.2.1.2, “Options File” on page 153 for further information.

We define a backup copy group for each of the management classes using the following commands:

```
adsm> define copygroup server server data type=Backup destination=DISKDATA \  
cont> frequency=1 verexists=3 verdeleted=1 retextra=100 retonly=365 \  
cont> mode=modified serialization=shrstatic  
ANR1530I Backup copy group STANDARD defined in policy domain server, set  
server, management class data.  
  
adsm> define copygroup server server directory type=Backup destination=DISKDIRS \  
cont> frequency=1 verexists=nolimit verdeleted=1 retextra=100 retonly=365 \  
cont> mode=modified serialization=shrstatic  
ANR1530I Backup copy group STANDARD defined in policy domain server, set  
server, management class directory.  
  
adsm> define copygroup server server special type=Backup destination=DISKDATA \  
cont> frequency=1 verexists=3 verdeleted=1 retextra=100 retonly=365 \  
cont> mode=modified serialization=shrdynamic  
ANR1530I Backup copy group STANDARD defined in policy domain server, set  
server, management class special.  
  
adsm> define copygroup workstn workstn data type=Backup \  
cont> destination=DISKDATA frequency=1 verexists=2 verdeleted=1 retextra=30 \  
cont> retonly=100 mode=modified serialization=shrstatic  
ANR1530I Backup copy group STANDARD defined in policy domain workstn, set  
workstn, management class data.  
  
adsm> define copygroup workstn workstn directory type=Backup \  
cont> destination=DISKDIRS frequency=1 verexists=nolimit verdeleted=1 retextra=30 \  
cont> retonly=100 mode=modified serialization=shrstatic  
ANR1530I Backup copy group STANDARD defined in policy domain workstn, set  
workstn, management class directory.  
  
adsm> define copygroup workstn workstn special type=Backup \  
cont> destination=DISKDATA frequency=1 verexists=2 verdeleted=1 retextra=30 \  
cont> retonly=100 mode=modified serialization=shrdynamic  
ANR1530I Backup copy group STANDARD defined in policy domain workstn, set  
workstn, management class special.
```

6.2.5 Defining the Archive Copy Group

For the recommended configuration, we define only archive copy groups for the default management classes. The data management class in the SERVER and WORKSTN policy domains keeps archive copies for 365 days, and does not archive files that are unavailable after a few retries.

```

adsm> define copygroup server server data type=Archive destination=DISKDATA \
cont> retver=365 serialization=shrstatic
ANR1535I Archive copy group STANDARD defined in policy domain server, set
server, management class data.

adsm> define copygroup workstn workstn data type=Archive \
cont> destination=DISKDATA retver=100 serialization=shrstatic
ANR1535I Archive copy group STANDARD defined in policy domain workstn, set
workstn, management class data.

```

6.2.6 Deleting the STANDARD Policy Domain

You should remove the STANDARD domain information to clean up your environment. The `delete domain` command removes an existing policy domain. This command is very powerful in that it also removes the policy sets, management classes, and copy groups that belong to the domain. For this reason, the `delete domain` command should be used with caution. The following command shows how to remove a policy domain named STANDARD.

```

adsm> delete domain standard

Do you wish to proceed? (Yes/No) yes
ANR1501I Policy domain STANDARD deleted.

```

If policy domain STANDARD contains nodes and filespaces, this command fails. You must remove all registered nodes from a policy domain before deleting it. There are two ways to accomplish this task:

- Delete the nodes and their filespaces
- Move the nodes including their filespaces to another domain

6.2.6.1 Deleting a Node

Deleting a node involves removing the node data with the `delete filesystem` command, and removing the node definition with the `remove node` command. The following commands show how to remove the filespaces and node definition for node `devnode1`.

```

adsm> delete filesystem devnode1 *
ANR2238W This command will result in the deletion of all inventory references
to the data on filespaces that match the pattern * for node DEVNODE1, whereby
rendering the data unrecoverable.

Do you wish to proceed? (Yes/No) yes
ANS8003I Process number 48 started.

adsm> remove node devnode1

Do you wish to proceed? (Yes/No) yes
ANR2061I Node DEVNODE1 removed from policy domain STANDARD.

```

6.2.6.2 Moving a Node

The `update node` command moves the node to a new domain. The data is moved to the new domain during the next backup operation. The process of moving the data to the new domain's management classes is called rebinding. If the new domain has the same management class names as the old one, very little change takes place, as the data is affected more by the management class name than the policy domain name. However, if the new domain has completely different management class names, then the data is rebound to the default management class for the domain. The first backup after the domain change may take longer than normal, but it is a one time operation and returns to normal speed thereafter.

Warning:

Care must be taken when migrating to another policy domain. If a client had data backed up under the old domain, the data is rebound to the same management class of the new domain, if it exists. If the management class does not exist in the new domain, the data is rebound to the default management class of the new domain. If this management class does not have a backup copy group, the data is rebound to the retention grace period defined in the new domain. This can lead to undesired expiration of backup versions.

The following command shows how to move the node `devnode1` from its current domain to the `WORKSTN` domain.

```
adsm> update node devnode1 domain=workstn
ANR2063I Node DEVNODE1 updated.

adsm> query node devnode1
```

Node Name	Platform	Policy Domain Name	Days Since Last Access	Days Since Password Set	Locked?
DEVNODE1	AIX	WORKSTN	1	1	No

6.3 Verifying Policy Definitions

The best way to check the policy definitions is to examine the details of the copy groups using the `query copygroup` command. The output from this command displays the policy domain, policy set, management class, and copy group names, as well as the copy group parameters. Note that it is really the copy group definitions that define the policy for the domain; the rest of the constructs between domain and copy group just provide flexibility in your configuration.

6.3.1 Backup Copy Groups

To check the settings for the recommended backup copy groups, use the `query copygroup` command to get the attributes displayed:

```

adsm> q copygroup server server type=backup

```

Policy Domain Name	Policy Set Name	Mgmt Class Name	Copy Group Name	Versions Data Exists	Versions Data Deleted	Retain Extra Versions	Retain Only Version
server	server	data	STANDARD	3	1	100	365
server	server	directory	STANDARD	3	1	100	365
server	server	special	STANDARD	3	1	100	365

```

adsm> q copygroup workstn workstn type=backup

```

Policy Domain Name	Policy Set Name	Mgmt Class Name	Copy Group Name	Versions Data Exists	Versions Data Deleted	Retain Extra Versions	Retain Only Version
workstn	workstn	data	STANDARD	2	1	30	100
workstn	workstn	directory	STANDARD	2	1	30	100
workstn	workstn	special	STANDARD	2	1	30	100

6.3.2 Archive Copy Groups

To check the settings for the recommended archive copy groups, use the query copygroup command to get a quick look at the attributes:

```

adsm> q copygroup server server type=archive

```

Policy Domain Name	Policy Set Name	Mgmt Class Name	Copy Group Name	Retain Version
server	server	data	STANDARD	365

```

adsm> q copygroup workstn workstn type=archive

```

Policy Domain Name	Policy Set Name	Mgmt Class Name	Copy Group Name	Retain Version
workstn	workstn	data	STANDARD	100

6.4 Validating and Activating a Policy Set

The last step in setting up your policy is to validate and activate your policy set. The commands are very straightforward and have few parameters.

6.4.1 Validating the Recommended Policy Sets

The validate policysset command checks for completeness in the management class and copy group definitions. It validates these policies and makes them ready for activation.

To validate the recommended policy sets:

```

adsm> validate policysset server server
ANR1515I Policy set server validated in domain server (ready for activation).

adsm> validate policysset workstn workstn
ANR1515I Policy set workstn validated in domain workstn (ready for activation).

```

6.4.2 Activating the Recommended Policy Sets

The `activate policyset` command allows the specified policy set to be applied to the data within its policy domain.

The following example shows how to activate the recommended policy sets:

```
adsm> activate policyset server server
Do you wish to proceed? (Yes/No) y
ANR1514I Policy set server activated in policy domain server.

adsm> activate policyset workstn workstn
Do you wish to proceed? (Yes/No) y
ANR1514I Policy set workstn activated in policy domain workstn.

adsm> query domain
```

Policy Domain Name	Activated Policy Set	Activated Default Mgmt Class	Number of Registered Nodes	Description
server	server	data	1	Servers Policy Domain
workstn	workstn	data	1	Workstations Policy Domain

6.5 Enforcing Your Policy

Once the policy is defined, you want to enforce it using the `expire inventory` command. This command makes sure that any extra copies of data in your copy groups are removed from the database. It also takes care of data that is older than your specified retention period. The backup copy group parameters `VEREXISTS`, `VERDELETED`, `RETEXTRA`, and `RETONLY` are applied to backup data through the expiration process as well as the archive copy group parameter `RETVR`. The following command shows how to expire the database references for the ADSM server.

```
adsm> expire inventory
Session established with server PALANA: AIX-RS/6000
Server Version 3, Release 1, Level 2.15
Server date/time: 03/03/1999 16:28:41 Last access: 03/03/1999 14:25:57
ANS8003I Process number 50 started.
```

This command can be computationally intensive; thus, it should only be executed when critical processes are completed.

The `EXPINTERVAL` parameter in the server options file specifies the number of hours between automatic expiration processing, and is initially configured for 24 hours. With automatic expiration processing enabled, the server runs inventory expiration at start-up and every 24 hours thereafter. We recommend setting the `EXPINTERVAL` to zero (disabling automatic expiration) and defining a daily administrative schedule to run this command at a convenient time.

Chapter 7. User Management

In this chapter, we explain the creation and maintenance of users of an ADSM system. There are two categories of users:

- Administrators
- Client nodes

Administrators manage ADSM resources. They also administer ADSM client nodes. One way of maintaining client nodes is to centrally define options the clients will use during backup, restore, archive, and retrieve operations. This is done using client option sets.

7.1 Administrators

An ADSM administrator manages ADSM resources. The number of administrators and their level of privileges will vary according to your environment. Administrators with system privilege can perform any ADSM function. Administrators with policy, storage, operator, or analyst privileges can perform subsets of ADSM functions. All commands issued by administrators are logged to the server activity log. See 2.6, “Administrator IDs” on page 32 for planning considerations.

7.1.1 Considerations

The creation of an administrator is a two-step process:

1. Define the administrator ID.
2. Grant the necessary privileges.

The `register admin` command defines administrators. The only required parameters are the admin name, or user ID, and initial password. You should consider using the optional `CONTACT` parameter to distinguish administrators.

The `grant authority` command grants the necessary privileges to an administrator. The only required parameters are the admin name and the privileges being granted.

The `query admin` command displays information about one or more administrators.

7.1.2 Default Environment

During the initial server startup, two administrators are automatically defined: `SERVER_CONSOLE` and `ADMIN`.

`SERVER_CONSOLE` is a special administrator associated with the server console that has system privileges. You cannot update, lock, rename, or remove the `SERVER_CONSOLE` user ID from ADSM. The `SERVER_CONSOLE` user ID does not have a password. Therefore, you cannot use the user ID from an administrative client unless you set authentication off. Although you can change the authority of this administrator, we recommend that you do not. In an emergency situation, administrative commands can be issued from the server console to correct situations such as forgotten administrator passwords and locked IDs.

ADMIN is an administrator that has system privileges and an initial password of ADMIN. This administrator is used to set up your ADSM environment. For enhanced security after your environment has been created, we recommend that you delete this administrator.

7.1.3 Recommended Administrators

We recommend that you define the following administrators to support your ADSM environment:

- System
- Support
- Reporting
- Client

We recommend that you use the optional CONTACT parameter to distinguish administrators. Although we have recommended certain administrator names, you do not have to use them. You can use any name suitable for your purposes.

In 1.2, “Redbook Support Material” on page 2, we show how to load a predefined macro into ADSM. The macro `mac.admins`, which we provide to define administrative IDs for our redbook environment, is shown in Appendix B.1.1, “Define Administrators” on page 243.

7.1.3.1 System

Define an ID for the ADSM administrator and possibly another for your backup ADSM administrator. These administrators should have system privileges.

In our example, we define an administrator named SYSADMIN, and a backup named SYSADMIN2. The initial password for both administrators is AMANDA. The administrative commands to create our recommended administrators, and the results of those commands, should look like this:

```
adsm> register admin sysadmin amanda contact='ADSM Administrator'  
ANR2068I Administrator SYSADMIN registered.  
  
adsm> grant authority sysadmin classes=system  
ANR2076I System privilege granted to administrator SYSADMIN.  
  
adsm> register admin sysadmin2 amanda contact='ADSM Administrator (Alternate)'  
ANR2068I Administrator SYSADMIN2 registered.  
  
adsm> grant authority sysadmin2 classes=system  
ANR2076I System privilege granted to administrator SYSADMIN2.
```

7.1.3.2 Support

Define an ID for your ADSM technical support person. This ID should have system privilege.

In our example, we define a support administrator named SUPPORT. The initial password for this administrator is KIRSTEN. The administrative commands to create our recommended administrator and the results of those commands should look like this:

```
adsm> register admin support kirsten contact='ADSM Support'  
ANR2068I Administrator SUPPORT registered.  
  
adsm> grant authority support classes=system  
ANR2076I System privilege granted to administrator SUPPORT.
```

7.1.3.3 Reporting

Define an ID for ADSM reporting purposes. Although this administrator has no special privileges, it does have authority to issue QUERY commands. This is especially useful for producing regular reports through a method such as scripts. Disclosure of this administrator's password does not create a security exposure as the administrator cannot change any ADSM resource.

In our example, we define a support administrator named REPORTER. The initial password for this administrator is CAMERON. The administrative command to create our recommended administrator and the results of that command should look like this:

```
adsm> register admin reporter cameron contact='ADSM Reporting'  
ANR2068I Administrator REPORTER registered.
```

7.1.3.4 Client

Define an ID for the Web administrative backup-archive client access. This administrator would have client access authority to all nodes in the ADSM environment. In particular, this administrator can perform restore operations on behalf of users. This is usually sufficient for smaller organizations. In larger organizations, you might require multiple administrators with this authority.

In our example, we define a client administrator named HELPDESK. The initial password for this administrator is PHIL. The administrative commands to create our recommended administrator and the results of those commands should look like this:

```
adsm> register admin helpdesk phil contact='ADSM Client Administrator'  
ANR2068I Administrator HELPDESK registered.  
  
adsm> grant authority helpdesk classes=node node=*  
ANR2126I GRANT AUTHORITY: Administrator HELPDESK was granted ACCESS authority  
for client PAGOPAGO.  
ANR2126I GRANT AUTHORITY: Administrator HELPDESK was granted ACCESS authority  
for client COCOS.  
ANR2126I GRANT AUTHORITY: Administrator HELPDESK was granted ACCESS authority  
for client PUTINEY.  
ANR2126I GRANT AUTHORITY: Administrator HELPDESK was granted ACCESS authority  
for client FIJI.
```

The administrator is only granted authority to those nodes that exist at the time the command is issued. If you create additional nodes and you want administrators to have authority to those nodes, the `grant authority` command must be reissued.

7.1.4 Working with Administrators

In this section we explain how to perform various processes related to administrators that you will find useful in your ADSM environment:

7.1.4.1 Displaying Administrators

You use the `query admin` command to display administrators. If no administrator name is specified, all administrators are displayed. If an administrator name is specified, only that administrator's information is displayed.

The administrative command to display the list of all administrators and the results of that command should look similar to this:

```
adsm> query admin
```

Administrator Name	Days Since Last Access	Days Since Password Set	Locked?	Privilege Classes
ADMIN	<1	12	No	System
HELPDESK	<1	<1	No	Client Access
REPORTER	<1	<1	No	
SERVER_CONSOLE			No	System
SUPPORT	<1	<1	No	System
SYSADMIN	<1	<1	No	System
SYSADMIN2	<1	<1	No	System

The administrative command to display full details of our administrator named HELPDESK and the results of that command should look similar to this:

```
adsm> query admin helpdesk format=detailed
```

```
Administrator Name: HELPDESK
Last Access Date/Time: 01-02-1999 15:37:36
Days Since Last Access: 1
Password Set Date/Time: 01-02-1999 15:37:36
Days Since Password Set: 1
Invalid Sign-on Count: 0
Locked?: No
Contact: ADSM Client Administrator
System Privilege:
Policy Privilege:
Storage Privilege:
Analyst Privilege:
Operator Privilege:
Client Access Privilege: COCOS PUTNEY PAGOPAGO FIJI
Client Owner Privilege:
Registration Date/Time: 01-02-1999 15:37:36
Registering Administrator: ADMIN
Managing profile:
```

7.1.4.2 Changing an Administrator Password

An administrative user can change his or her own password using the `update admin` command. A user with system privilege can update any administrator's password. There is no way of determining the current password for an administrator.

The administrative command to change the password for the administrator named SYSADMIN2 to DAVID should look like this:

```
adsm> update admin sysadmin2 david
ANR2071I Administrator SYSADMIN2 updated.
```

When this command is logged to the server activity log, the password in the log is replaced by asterisks.

7.1.4.3 Deleting the ADMIN Administrator

After you have defined the set of system administrators, delete the installed system administration ID, ADMIN. As you cannot delete an administrator if that administrator is currently accessing the server, you must logon to ADSM as a different administrator to perform this task. The administrative session, the administrative command, and the results of that command should look similar to this:

```
C:\Program Files\IBM\adsm\saclient>dsmadmc

ADSTAR Distributed Storage Manager
Command Line Administrative Interface - Version 3, Release 1, Level 0.6
(C) Copyright IBM Corporation, 1990, 1997, All Rights Reserved.

Enter your user id: sysadmin

Enter your password: *****

Session established with server PAGOPAGO: Windows NT
  Server Version 3, Release 1, Level 2.13
  Server date/time: 01-02-1999 16:08:32 Last access: 01-02-1999 15:36:12

adsm> remove admin admin

Do you wish to proceed? (Yes/No) yes
ANR2069I Administrator ADMIN removed.
```

7.2 Client Nodes

An ADSM client must be registered with the ADSM server before any backup and recovery operations can be performed for that client. Although ADSM provides two modes, open and closed, for registering client nodes to the server, we recommend that you do not change the closed mode default.

7.2.1 Considerations

The registration of a client node is a single-step process using the `register node` command. Registration requires an ADSM node name and a client access password as a minimum. You should use the machine name for the ADSM node name. If you define your own ADSM policy domains, then you must also use the `parameter`. Further, we recommend that you also specify the `USERID` parameter.

The `DOMAIN` parameter specifies the name of the policy domain to which the node is assigned. If you do not specify a policy domain name, the node is assigned to the default policy domain (STANDARD). Within our recommended environment, there are two domains, `SERVER` and `WORKSTN`, and no `STANDARD` domain. Consequently, if you use our environment, you must specify the `DOMAIN` parameter when registering a client node.

The USERID parameter specifies the name of an ID who will be defined as an administrator with node owner authority to the client. Node authority allows an administrator to use the Web backup-archive client. When the USERID parameter is omitted, an administrator with the same name as the node is defined. This is the default. We recommend that you specify USERID=NONE for our environment and grant node authority to the HELPDESK administrator and other administrators explicitly through the `grant authority` command.

7.2.2 Default Environment

Except for MVS, one client node is registered with the ADSM server during the initial server startup. This client node is the ADSM server itself. The name of this client is CLIENT. The initial password is CLIENT. The node is assigned to the STANDARD domain.

7.2.3 Working with Client Nodes

In this section, we explain how to perform various processes related to client nodes that you will find useful in your ADSM environment.

7.2.3.1 Registering a Client Node

You use the `register node` command to define a client to the server. In our environment, there are parameters (node name, node password, domain name) which you must supply, and two parameters (contact, userid) which you should supply to register a client node.

Table 29 on page 128 shows an extracted version of Table 7., “Client Requirements Worksheet” on page 9 completed with a set of sample client information.

Table 29. Client Requirements Worksheet: Example

Node name	Contact	Domain
PAGOPAGO	NTA Group	SERVER
COCOS	phil@home	WORKSTN
FIJI	FlightDeck	SERVER
PUTNEY	Ops x222	SERVER

Assuming a common password of KOALA, the commands to define this set of client nodes and resulting output should look like this:

```
adsm> register node pagopago koala domain=server userid=none contact='NTA Group'
ANR2060I Node PAGOPAGO registered in policy domain SERVER.

adsm> register node cocos koala domain=workstn userid=none contact='phil@home'
ANR2060I Node COCOS registered in policy domain WORKSTN.

adsm> register node fiji koala domain=server userid=none contact=FlightDeck
ANR2060I Node FIJI registered in policy domain SERVER.

adsm> register node putney koala domain=server userid=none contact='Ops x222'
ANR2060I Node PUTNEY registered in policy domain SERVER.
```

7.2.3.2 Displaying Client Nodes

You use the `query node` command to display client nodes. If you do not specify a node name, all nodes are displayed. If you do specify a node name, only that node's information is displayed.

The administrative command to display the list of all nodes and the results of the command should look similar to this:

```
adsm> query node
```

Node Name	Platform	Policy Domain Name	Days Since Last Access	Days Since Password Set	Locked?
COCOS	WinNT	WORKSTN	<1	<1	No
FIJI	AIX	SERVER	<1	<1	No
PAGOPAGO	WinNT	SERVER	4	12	No
PUTINEY	AIX	SERVER	<1	<1	No

The administrative command to display full details of the node named `PAGOPAGO`, and the results of that command, should look similar to this:

```
adsm> query node pagopago format=detail
```

```
Node Name: PAGOPAGO
Platform: WinNT
Client OS Level: 4.00
Client Version: Version 3, Release 1, Level 0.6
Policy Domain Name: SERVER
Last Access Date/Time: 29-01-1999 13:56:23
Days Since Last Access: 4
Password Set Date/Time: 21-01-1999 17:55:29
Days Since Password Set: 12
Invalid Sign-on Count: 0
Locked?: No
Contact: NIA Group
Compression: Client's Choice
Archive Delete Allowed?: Yes
Backup Delete Allowed?: No
Registration Date/Time: 20-01-1999 17:15:33
Registering Administrator: ADMIN
Last Communication Method Used: Named Pipe
Bytes Received Last Session: 15,680
Bytes Sent Last Session: 17,038
Duration of Last Session (sec): 8,677.00
Pct. Idle Wait Last Session: 4.85
Pct. Comm. Wait Last Session: 0.00
Pct. Media Wait Last Session: 0.00
Optionset:
URL:
Node Type: Client
```

7.2.3.3 Granting Access to Client Nodes

An administrator can perform all client operations on behalf of the user through the Web backup-archive client. The administrator needs node authority to perform those operations. An administrative user with system privilege can grant that authority.

The administrative command to grant node authority to a client administrator named `HELPDESK`, and the results of that command, should look like this:

```

adsm> grant authority helpdesk classes=node node=*
ANR2126I GRANT AUTHORITY: Administrator HELPDESK was granted ACCESS authority
for client COCOS.
ANR2126I GRANT AUTHORITY: Administrator HELPDESK was granted ACCESS authority
for client FLJI.
ANR2126I GRANT AUTHORITY: Administrator HELPDESK was granted ACCESS authority
for client PAGOPAGO.
ANR2126I GRANT AUTHORITY: Administrator HELPDESK was granted ACCESS authority
for client PUTNEY.

```

7.2.3.4 Changing Passwords for Client Nodes

An administrative user with either system privilege, unrestricted policy privilege, or restricted policy privilege can change the password of a client node.

The administrative command to change the password for the client node named COCOS to ISLAND and the results of that command should look like this:

```

adsm> update node cocos island
ANR2063I Node COCOS updated.

```

7.2.3.5 Deleting Filespaces for a Client Node

An administrative user with either system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the client node is assigned, can delete node filesystems. Filespaces contain a client's backup and archive data. There are one or more filesystems for each client.

The `delete filesystem` command creates a server process that deletes one or more file spaces as a series of batch database transactions, thus preventing a rollback or commit for an entire file space as a single action. If the `delete filesystem` process is canceled or if a system failure occurs, a partial deletion can occur. In this case, a subsequent `delete filesystem` command for the same node can delete the remaining data. You must delete all backup and archive file spaces that belong to a client node before you can delete that client node.

The administrative command to query the filesystems that belong to the client named COCOS, and the results of that command, should look similar to this:

```

adsm> query filesystem cocos

```

Node Name	Filespace Name	Platform	Filespace Type	Capacity (MB)	Pct Util
COCOS	\\cocos\c\$	WinNT	FAT	1,612.2	58.7
COCOS	\\cocos\d\$	WinNT	NIFS	1,612.4	0.7

The administrative command to delete the filesystem named \\cocos\c\$ for the client node named COCOS, and the results of that command, should look similar to this:

```
adsm> delete filespace cocos \\cocos\c$
ANR2238W This command will result in the deletion of all inventory references
to the data on filespace that match the pattern \\cocos\c$ for node COCOS,
whereby rendering the data unrecoverable.

Do you wish to proceed? (Yes/No) yes
ANS8003I Process number 6 started.
```

Filespace deletions can run as either a background or foreground process. The previous example was run as a background process. You must monitor the completion of that process through the `query process` command. You must view messages from that process on the server activity log.

You can run filespace deletion in the foreground by adding the `WAIT=YES` parameter to the command. Foreground processes write messages to your session. If your filespace contains a large number of files, we recommend that you run the process in the background.

The administrative command to delete all filespace for the client node named COCOS in the foreground, and the results of that command, should look similar to this:

```
adsm> delete filespace cocos * wait=yes
ANR2238W This command will result in the deletion of all inventory references
to the data on filespace that match the pattern * for node COCOS, whereby
rendering the data unrecoverable.

Do you wish to proceed? (Yes/No) yes
ANR0984I Process 8 for DELETE FILESPACE started in the FOREGROUND at 14:19:20.
ANR0800I DELETE FILESPACE: * for node COCOS started as process 4.
ANR0802I Delete Filespace * (backup/archive data) for node COCOS started.
ANR0806I Delete Filespace * complete for node COCOS: 19 files deleted.
ANR0987I Process 4 for DELETE FILESPACE running in the FOREGROUND processed 19
items with a completion state of SUCCESS at 14:19:20.
```

7.2.3.6 Deleting a Client Node

An administrative user with either system privilege, unrestricted policy privilege, or restricted policy privilege can delete a client node. Before you can remove a client node, you must delete all backup and archive file spaces that belong to that client node.

The administrative command to delete the client node named COCOS, and the results of that command, should look similar to this:

```
adsm> remove node cocos

Do you wish to proceed? (Yes/No) yes
ANR2061I Node COCOS removed from policy domain WORKSTIN.
```

7.3 Client Option Sets

An ADSM client session has a set of options which are used during the backup, archive, restore, or retrieve processes. Options can be specified in two ways:

1. In the client options file. This is mandatory during the set up of a client. A client options file is a set of ADSM client options stored in one or two (UNIX clients only) files on the client.
2. Using the client options set. This is optional. A client option set is a set of ADSM client options stored in the ADSM database. An option set can be associated with one or more ADSM clients, but an ADSM client can be associated with only one option set.

7.3.1 Considerations

The options defined in a client option set are a subset of the available client options. Options such as communications are still stored on the client machine. When the same individual option is specified in both the local options file and the options set, the default is that the options file version is used. However, you can specify that individual options in an option set cannot be overridden in the client's local option file. Although include-exclude specifications cannot be overridden, you can specify the sequence in which the option set specifications are processed. Thus one set of default values can be defined for each type of client, and the client machines can still be customized, within acceptable limits.

We recommend that you use client option sets for ease of administration. Management of the environment is complex where the number of clients is growing and the number of options is increasing. The use of client option sets eases that administrative burden by centralizing the management of those options and clients. It is easier to update a client options set once, than to perform the same update to the local client options file on each node.

When you specify include-exclude options in both an options file and an options set, understanding how they relate to each other is critical to understanding what actually happens. The sequence is not obvious.

Include-exclude options in the client options set are additive options and will not override the include-exclude options in the client options file. The sequence number in the client option set determines the order in which the include-exclude statements are added to the existing include-exclude statements of the client options file. The statements from the client options set are added at the end of the include-exclude statements of the options file, in sequence order.

For example, suppose you have the following specifications in a client option set:

```
include c:\test\* seq=1
exclude c:\working\* seq=2
```

and the following exclude statement in your options file:

```
exclude c:\test\*
include c:\working\*
```

The complete set of include-exclude specifications would be:

```
exclude c:\test\*
include c:\working\*
include c:\test\*
exclude c:\working\*
```

Note:

Include-exclude statements are read and processed from the bottom to the top.

Processing using this set of include-exclude statements would result in the files under `C:\working` not being backed up, and those files under `C:\test` being backed up.

7.3.2 Default Environment

There are no client option sets created by default in an installed ADSM environment.

7.3.3 Recommended Client Option Sets

In 1.2, “Redbook Support Material” on page 2, we show how to load a predefined macro into ADSM. The macro `mac.optionsets`, which we provide to define client option sets for our redbook environment, is shown in Appendix B.1.2, “Define Client Option Sets” on page 244.

In our environment, we use multiple client option sets. All option sets are built from the base option set named `redbook`. This option set contains options other than the include-exclude specifications. All other option sets are based on client platform. Each is built from the base option set with the addition of specific include-exclude recommendations for that platform.

7.3.4 Associating a Client Node with a Client Option Set

The client node definition should be updated to use a client option set. You need either system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the client node belongs to issue the required `update node` command.

You can use the `select` command to view the names of all the client option sets. Do not use the `query` command, as it generates too much output, and you cannot easily distinguish the client option set names.

The administrative command to show names of all client option sets and the results of that command should look similar to this:

```
adsm> select * from cloptsets
```

OPTIONSET_NAME	DESCRIPTION	LAST_UPDATE_BY	PROFILE
AIX	AIX Clients	SYSADMIN	
NETWARE	Netware Clients	SYSADMIN	
REDBOOK	Redbook Base Set	SYSADMIN	
WINDOWS	Windows Clients	SYSADMIN	

The administrative command to associate a client option set named AIX with the client node named PUTNEY, and the results of that command, should look like this:

```
adsm> update node putney cloptset=aix
ANR2063I Node PUTNEY updated.
```

7.3.5 Working with Client Option Sets

In this section we explain how to perform various processes related to client option sets that you will find useful in your ADSM environment.

7.3.5.1 Creating a Client Option Set

Defining a client option set consists of two steps:

1. Creating the option set
2. Populating it with client options

You create a new client option set with the `define cloptset` command or you can clone an existing option set with the `copy cloptset` command. You can add an option to an option set with the `define clientopt` command. You can remove an option from an option set with the `delete clientopt` command. You need system privilege or unrestricted policy privilege to issue these commands.

The administrative commands to define a client option set named ALLNODES, populate it with options, and the results of those commands, should look like this:

```
adsm> define cloptset allnodes
ANR2046I DEFINE CLOPTSET: Optionset ALLNODES defined.

adsm> define clientopt allnodes compression yes
ANR2050I DEFINE CLIENTOPT: Option COMPRESSION defined in optionset ALLNODES.

adsm> define clientopt allnodes maxcmdretries 4
ANR2050I DEFINE CLIENTOPT: Option MAXCMDRETRIES defined in optionset ALLNODES.

adsm> define clientopt allnodes verbose ''
ANR2050I DEFINE CLIENTOPT: Option VERBOSE defined in optionset ALLNODES.
```

The administrative commands to clone the client option set named ALLNODES to another named CLONE, replace the verbose option with the quiet option, and the results of those commands, should look like this:

```
adsm> copy cloptset allnodes clone
ANR2055I COPY CLOPTSET: Optionset ALLNODES copied to optionset CLONE.

adsm> delete clientopt clone verbose
ANR2053I DELETE CLIENTOPT: Option VERBOSE, sequence number 0, has been deleted
from optionset CLONE.

adsm> define clientopt clone quiet ''
ANR2050I DEFINE CLIENTOPT: Option QUIET defined in optionset CLONE.
```

7.3.5.2 Deleting a Client Option Set

You need either system privilege or unrestricted policy privilege to delete a client option set with the `delete cloptset` command.

The administrative command to delete the option set named ALLNODES, and the results of that command, should look like this:

```
adsm> delete cloptset allnodes  
Do you wish to proceed? (Yes/No) yes  
ANR2048I DELETE CLOPTSET: Optionset ALLNODES deleted.
```

The association between a node and a client option set is removed automatically when that client option set is deleted. The node is not associated with another client option set until you reissue the `update node` command.

Chapter 8. Licensing Your ADSM System

This chapter describes the tasks of licensing an ADSM system and monitoring its compliance. We also show how to migrate from an ADSM evaluation (try-and-buy version) to a full ADSM version. See 2.7, “License Considerations” on page 33 for planning considerations.

8.1 Licensed Features

The base ADSM server license supports an unlimited number of administrative clients, one backup-archive client using named pipes protocol, and a selection of removable media devices. Table 30 on page 137 lists all currently available licensed features you can add on to your base ADSM server license.

Table 30. ADSM Licensed Features

Licensed feature	License files
Additional backup-archive clients: The base ADSM license supports one backup-archive client	1client.lic, 5client.lic 10client.lic 50client.lic
Network connections: To use a network communication method, for example, TCP/IP	network.lic
Open Systems Environment client support	opensys.lic
Hierarchical storage management (HSM) clients, also known as space-managed clients	spaceman.lic
Disaster Recovery Manager (DRM)	drm.lic
Enterprise Administration services	entadmin.lic
Server-to-server virtual volumes support	virtvols.lic
Advanced device support	advdev.lic

The enrollment certificate files for all ADSM licenses are on the ADSM installation CD-ROM. You register those licenses you want by issuing the `register license` command with the name of the enrollment certificate file. When registered, the licenses are stored in a file named `nodelock` in the current directory.

Working through a licensing example should help you understand which licenses you will need for your configuration.

Consider the following scenario. You wish to install your ADSM server on a Windows NT system. There are different client platforms that wish to use the ADSM service, namely, 60 AIX clients, 5 Sun Solaris clients, and 10 Windows clients. An IBM 3575 library L06 with 2 drives is your preferred choice for library use. Which licenses are required?

The easiest way to work through this is by following each license type.

- Base License: 1 x Windows NT single server license, one server install for the Windows NT platform. The single server edition includes a license for one backup-archive client through named pipes protocol, that is, on the same

machine as the server only, and for administrators connected to the server by any method.

- Network Connections for clients: 1 x Windows NT Network Enabler licenses all client types to communicate with the ADSM server over the network.
- Additional Backup-Archive Clients: There are a total of 75 clients, 76 if you include the ADSM server itself, which is covered by the single server license. The User registration licenses one backup-archive client. The client licenses you buy are for the type of server platform—not for the type of client. You need 1 x NT 50 User registration, 2 x NT 10 User registration and 1x NT 5 User registration.
- Advanced Device Support: The library you wish to use is an IBM 3575 library L06 with 2 drives. You need to find out whether this model is covered by the base supported devices list (no additional license required), or whether it will require the extended device support license. Check the URL: <http://www.storage.ibm.com/software/adsm/adsercli.htm#servfix> and click on the supported devices list for the Windows NT server platform. In our example, you need 1 x NT Extended Device Support for this model with 2 drives.

8.2 Registering Licensed Features

If you received an ADSM evaluation (try-and-buy), the license registration is done during the installation of the server.

If you bought a base ADSM server license, you can obtain licenses for licensed features, and register those licenses by specifying the license files using the `register license` command:

You can register any or all of these features. For example, if you want to add 10 additional backup-archive clients, issue the following command:

```
adsm> register license file(10client.lic)
ANR2852I Current license information:
ANR2835I Server is licensed for 67 clients.
ANR2861I Server is licensed to support NETWORK connections.
ANR2869I Server is licensed for Advanced Device Support.
ANR2853I New license information:
ANR2835I Server is licensed for 77 clients.
ANR2861I Server is licensed to support NETWORK connections.
ANR2869I Server is licensed for Advanced Device Support.
```

You can also register a license by specifying the product password that is included in the license certificate file.

8.3 Saving Your Licenses

When license registration is complete, the licenses are stored in a file named NODELOCK in the server start directory.

Save the CD-ROM containing your enrollment certificate files if you need to register your licenses again for any of the following reasons:

- The server is corrupted.
- The server has moved to a different machine.
- The NODELOCK file is destroyed or corrupted. ADSM stores license information in the NODELOCK file, which is located in the directory from which the server is started.

8.4 License Compliance

If license terms change, for example, if 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.

If the server uses a licensed feature but the license is not registered, the function fails. When you issue a command associated with an unlicensed feature, ADSM does not issue a warning message, and the command fails.

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.

In either case, you must contact your IBM account representative or authorized reseller to modify your agreement.

8.5 Monitoring Licenses

There are two commands to monitor the license registration on your ADSM system: `query license` and `audit license`.

8.5.1 Displaying License Information

Use the `query license` command to display details of your current licenses and determine licensing compliance.

```

adsm> query license

                Last License Audit: 03/18/1999 11:51:32
                Registered Client Nodes: 5
                Licensed Client Nodes: 77
                Are network connections in use ? : Yes
                Are network connections licensed ? : Yes
Are Open Systems Environment clients registered ? : No
  Are Open Systems Environment clients licensed ? : No
    Is space management in use ? : No
    Is space management licensed ? : No
  Is disaster recovery manager in use ? : No
  Is disaster recovery manager licensed ? : No
Are Enterprise Administration services in use ? : No
Are Enterprise Administration services licensed ? : No
  Are Server-to-Server Virtual Volumes in use ? : No
  Are Server-to-Server Virtual Volumes licensed ? : No
    Is Advanced Device Support required? : Yes
    Is Advanced Device Support licensed ? : Yes
                Server License Compliance: Valid

```

8.5.2 Auditing Licenses

An administrator can monitor license compliance by issuing the administrative command `audit licenses`. This command is used to compare the current configuration with the current licenses.

```

adsm> audit licenses
Session established with server PALANA: AIX-RS/6000
  Server Version 3, Release 1, Level 2.15
  Server date/time: 03/18/1999 16:29:14 Last access: 03/18/1999 16:04:24

ANR2817I AUDIT LICENSES: License audit started as process 31.
ANS8003I Process number 31 started.

```

Note:

During a license audit, the server calculates, by node, the amount of backup, archive, and space management storage in use. This calculation can take a great deal of CPU time and can stall other server activity. Use the `NOAUDITSTORAGE` server option to specify that storage is not to be calculated as part of a license audit.

8.5.3 Scheduling Automatic License Audits

Use the `set licenseauditperiod` command to specify the number of days between automatic audits performed by the ADSM server.

```

adsm> set licenseauditperiod 30
ANR2814I SET LICENSEAUDITPERIOD: License audit period changed to 30 days.

```

8.6 ADSM Evaluation License

A 60-day evaluation (try-and-buy version) of ADSM is available through the ADSM Business Partner or IBM representative. This try-and-buy version provides the same functions as the standard licensed product.

The only difference is that with the try-and-buy product, the licensing package is not available and installed. You are limited to 50 ADSM clients. The try-and-buy product differs slightly from the standard licensed product in both content and installation.

MVS Users:

There is no ADSM evaluation for MVS available.

You can apply maintenance to a try-and-buy version of ADSM.

If you later choose to buy ADSM, you can install the licensed product over the try-and-buy product. To upgrade to the licensed product, do the following:

- On Windows NT:
 1. Mount the ADSM licensed product CD in your CD-ROM drive.
 2. Start installation, choose customized installation, and select only the license files.
- On UNIX:
 1. Mount the ADSM licensed product CD in your CD-ROM drive.
 2. Install the license package (`adsm.license` on AIX).

Only the license file is installed, and no server or user data will be lost. After installing the license package, register any licensed features that you purchased as described in 8.2, “Registering Licensed Features” on page 138.

Chapter 9. Administrative Client

The ADSM administrative client provides an interface to an ADSM server that allows an ADSM administrator to manage and control ADSM resources. This chapter discusses how to install an ADSM administrative client and what options to configure. Some examples are shown of the interfaces for various types of access.

There are two types of administrative clients in ADSM Version 3:

- The administrative client code which includes the command line interface on most client platforms
- The Web administrative interface

Windows Users:

The ADSM Windows client also provides an ADSM administrative graphical user interface (GUI). However, since the introduction of the Web administrative interface, its usage less frequently required, and therefore it is not further mentioned in this redbook.

MVS Users:

A TSO administrative interface is available on MVS platforms, but is not discussed here. See the *ADSM V3R1 MVS Quick Start*, (GC35-0276) for more information.

Confusion runs rampant in trying to differentiate between the Web administrative interface and the Web backup-archive client. These two totally separate functions have similar names, are accessed by similar means, and both use IDs that are defined as administrators.

The Web administrative interface connects to an ADSM server via the Web to present an interface to a true ADSM administrator session. It is used by an ADSM administrator to configure, monitor, and control an ADSM server just as they would use the command line interface.

The Web backup-archive client connects to an ADSM backup-archive client via the Web to present an interface that remotely performs the functions of the backup-archive client on that client. It is used by ADSM administrators, help desk personnel, or end users to manually back up, restore, archive, or retrieve files on a remote client without having to physically visit the client. To enforce security, the Web backup-archive client connects to a backup-archive client using an ADSM administrator ID with NODE authority. This eliminates unique security definitions on every client for the ADSM administrators, help desk personnel, or end users. For more information about the Web backup-archive client, see 10.5, "Web Client Usage" on page 179.

9.1 Code Installation

ADSM server and client code fixes and enhancements are released on a regular basis. The fixes are available from IBM via the internet or on CD-ROM. See 3.1, “Latest Code Updates” on page 37 for further information.

This section describes the installation procedures for the following:

- Administrative command line interface
- Web administrative interface

Use of any administrative client requires defining an administrator ID with the appropriate authority. See 7.1.3, “Recommended Administrators” on page 124 for information on creating an administrator ID.

9.1.1 Administrative Client Code

Typically, the administrative client code is installed on a remote workstation. Installation of the ADSM administrative client code is an option that can be selected during the installation of the backup-archive client code. For information on installing your version of the client software, see *ADSM V3R1 Installing the Clients* (SH26-4080).

9.1.2 Web Administrative Interface

Installation of the Web administrative interface happens automatically during the installation of the ADSM server code. To enable the Web administrative interface, place the following options in the server option file `dsmserv.opt` on the ADSM server:

- `COMMETHOD HTTP` to select the communication method
- `HTTTPORT 1580`, or let them default, to select the TCP/IP port that ADSM will use to communicate with the Web administrative interface

Note:

At server initialization, the server reads the server options file. If you update a server option by editing the file, you must remember to stop and restart the server to activate the updated server options file.

The Web administrative interface requires a Java 1.1.5-capable Web browser such as Microsoft Internet Explorer 4.01 or Netscape 4.06 on the administrator’s workstation. Using a Web browser means that ADSM code is not required on the Web browser workstation.

9.2 Customization

The following describes the various settings for the administrative interfaces.

9.2.1 Administrative Client

Use of the command line interface requires a client options file to be present on the client.

Client options for an administrative client (other than the Web administrative interface) are set in the client option files `dsm.opt`. Assuming a TCP/IP communications method, the only required option is the `TCPSERVERNAME`. We strongly recommend using a DNS name instead of the dotted IP address (`xxx.xxx.xxx.xxx`) for the `TCPSERVERNAME`. A dotted IP address is inflexible, because it creates management issues if the ADSM server is moved, or the addressing scheme is changed. Other useful parameters are the setting the date, time, number and language format options. Table 31 on page 145 summarizes some of the administrative client options.

Table 31. Administrative Client Options

	Option	Notes
<code>TCPSERVERNAME</code>	kindu	The DNS name or IP address
<code>TCPPORT</code>	1500	Default
<code>DATEFORMAT</code>	3	YYYY-MM-DD
<code>TIMEFORMAT</code>	1	hh:mm:ss
<code>LANGUAGE</code>	AMENG	American English
<code>NUMBERFORMAT</code>	1	1,000.00

9.2.2 Web Administrative Interface

No setup is required to use the Web administrator client, assuming that you have a compatible Web browser and TCP/IP available on the administrative workstation.

9.3 Interfaces

This section discusses the command line and the Web administrative interfaces available with ADSM, and also the Web backup-archive client as an administrative interface for the backup-archive client.

9.3.1 Command Line Interface

The command line interface has the benefit of being a character mode interface, and thus is well suited for those users who want to type the commands. You may also consider using it when you cannot access the Web administrative interface.

To start the command line interface, you must execute the administrative client program by typing `DSMADMC` in the operating system prompt. You are then prompted for an administrative user ID and password. Without authenticating, you cannot logon to the ADSM server. The following shows the start of an administrative client session on UNIX:

```

kindu:[/usr/lpp/adsm/bin]$ dsmdmc
ADSTAR Distributed Storage Manager
Command Line Administrative Interface - Version 3, Release 1, Level 0.6
(C) Copyright IBM Corporation, 1990, 1997, All Rights Reserved.

Enter your user id: sysadmin

Enter your password:

Session established with server PAGOPAGO: Windows NT
Server Version 3, Release 1, Level 2.13
Server date/time: 02/11/1999 11:23:09 Last access: 02/11/1999 11:04:49

adsm>

```

You then get the prompt `adsm>` to type the administrative ADSM commands. Once an administrative command is executed, you always return back to the `adsm>` prompt. This interactive method is also called *loop* mode. If you use the interactive mode, you do not have to re-enter your password with each command.

Another method to start an administrative session is the so-called *batch* mode. You simply type DSMADMC together with the administrative ID, the password, and the actual command that you want to execute. In this mode, the administrative client processes the command and returns to the operating system prompt.

Help is available from the `adsm>` command prompt by typing `help xxx` where `xxx` is blank, a command, or an ADSM message number.

A variant of the command line interface is available from inside the server utilities supplied with the Windows server code.

9.3.2 Web Administrative Interface

In ADSM Version 3 the Web administrative interface is the primary interface for administration. It incorporates Java applets and provides the following functions:

- ADSM enterprise console

The enterprise console displays ADSM as a single entity in the enterprise. It is an interface that integrates ADSM server and client functions for the administrator as a single application to manage ADSM in a distributed environment

- Server function

The Web administrative interface is an ADSM server function and consists of an integrated Web server and a new HTTP server communications protocol.

- Browser requirements

The Web administrative interface can be used from any workstation running an HTML 3.0 compliant Web browser with support for Java 1.1.5 or higher. Netscape 4.03, 4.04, and 4.05 requires a JDK 1.1.5 upgrade for this support. Netscape 4.06 has the required Java support. Microsoft's Internet Explorer 4.01 has the required Java support.

- Security

Security for the Web administrative interface is based on enhanced logging of Web-based sessions, password restrictions, and administrator lockout capabilities. In addition, Secure Sockets Layer (SSL) communications has been introduced in Version 3.1.2.

- Administrative functions

The Web administrative interface can be used to perform all administrative functions graphically, and it provides a command line and server event viewer.

The Web administrative interface connects to an ADSM server through the Web to present an interface to a true ADSM administrator session. Using a Web browser means that you actually do not have to have any ADSM code on the Web browser workstation.

To access the Web administrator interface, start your Web browser. In the Location or Address field, enter the URL:

`http://<servername>:1580`

where <servername> is the TCP/IP DNS name or dotted IP address of your ADSM server machine, and 1580 is the TCP/IP port number set by the HTTPPORT option.

Figure 19 on page 147 shows the Web administrative interface logon screen.

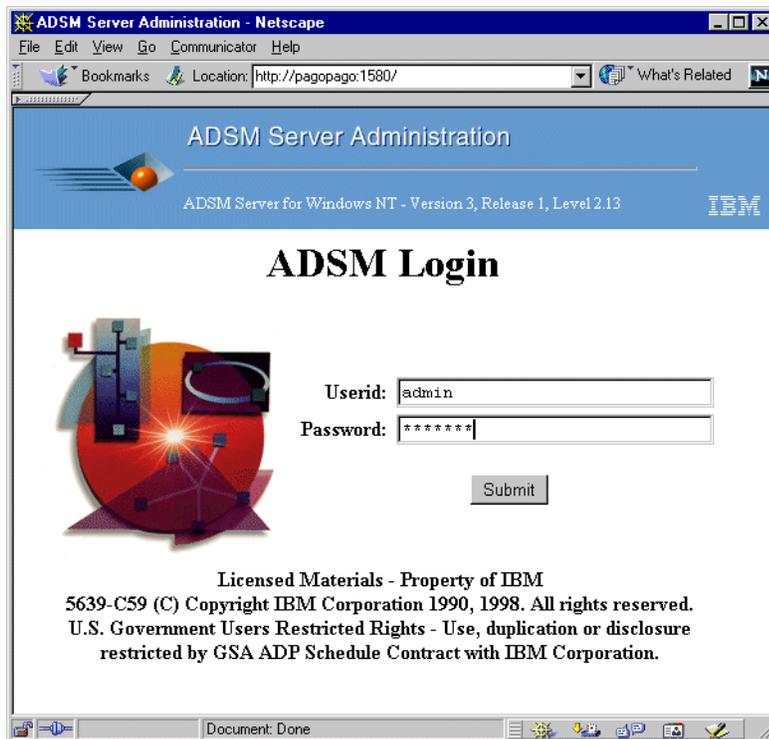


Figure 19. Web Administrative Interface Logon Screen

When an administrator logs in to the server through the Web administrative interface, the ADSM enterprise console is displayed. The enterprise console consists of three frames. Additional command line and event viewers can be displayed.

- **Banner Frame:** The banner frame indicates which server you are connected to (server name, platform, and version, release, level) and the fact that you are using the ADSM enterprise console. It also indicates the name you are logged on to as administrator.

The only operable control on the banner frame is a selection button that you can use to choose whether you want to display the command line and the event viewer. You can turn those windows off or on, depending on what you want to see in your browser display. Another option offers you the ability to log off (forcing reauthentication).

- **Tree Frame:** A collapsible tree is used for navigation. The content and layout are determined by the “view” that you have chosen. The initial branches of the tree are the views from which you can choose. Views are simply paradigms that are used to navigate to the objects in the enterprise that you want to use. Available views are operation, network, configuration, and object.
- **Detail Frame:** The detail frame displays detailed information about the item that you have selected from the tree frame. All operations that are available to manipulate the selected item are provided in the form of a pull-down menu in the top-right corner of the detail frame. For example, when a server group is selected, its attributes are displayed in the detail frame. Operations are provided for adding servers to the group, removing servers from the group, and defining new server groups.
- **Command Line:** An administrative command line can be displayed by selecting it from the selections pull-down menu in the banner frame.
- **Event Viewer:** A server event viewer can be displayed by selecting it from the selections pull-down menu in the banner frame.

Figure 20 on page 149 shows the ADSM enterprise console with the Server Status screen and an open drop-down box with the Set password expiration option selected. Also shown are the administrative command line and the event viewer.

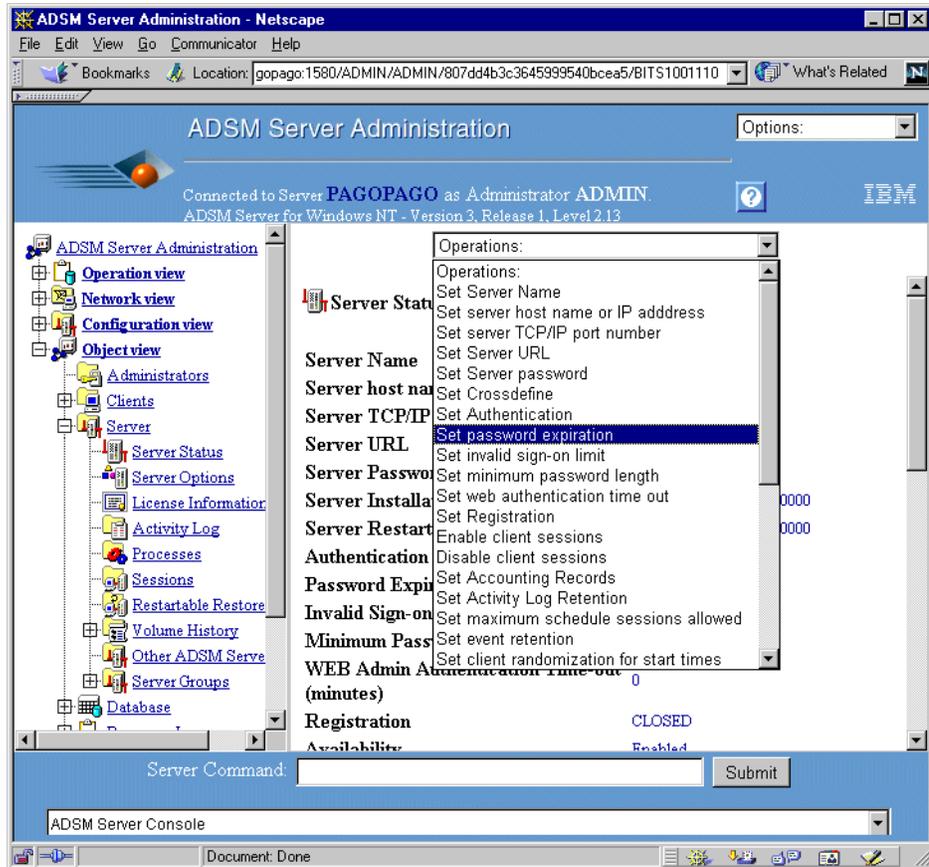


Figure 20. ADSM Enterprise Console in Web Administrative Interface

The Web administrative interface has a drop-down box in the upper right corner of the screen that allows you to perform three functions: show a command line, show an event line, and log off the ADSM administrator. In the command line area, you can issue any ADSM command that could be issued from the command line interface. The event line displays the command issued as a result of performing a function using the Web administrative interface. It maintains a history of the commands that have been issued in this session. Both the command line and the event line can be displayed at the same time.

The Web administrator session must re-authenticate after the time set in the ADSM server setting WEBAUTHTIMEOUT expires. At the next interaction with the ADSM server, you are prompted to enter your administrator ID and password. This time frame is based on wall clock time, not on the time since the last transaction. For example, if WEBAUTHTIMEOUT 10 is set, then every 10 minutes you will be prompted to re-authenticate, whether you have been active in those 10 minutes or not. Setting WEBAUTHTIMEOUT to zero indicates no timeout will occur.

We recommend that you use the Web administrative interface because it is platform-independent, it is location-independent (you only need a Web browser) and it allows you to perform all ADSM administrator functions.

Chapter 10. Backup-Archive Client

This chapter covers the steps you need to perform after installing the code from the installation media. It shows you how to use the backup-archive client interactively by using the command line interface and the GUI. It also shows you how to configure the backup-archive client to work either as a foreground program or as a background process, so that you can automate the backup processes using the scheduler facility. See 2.1, “Client Environment Data” on page 9 for planning considerations.

10.1 Code Installation

ADSM server and client code fixes and enhancements are released on a regular basis. The fixes are available from IBM through the internet or on CD-ROM. See 3.1, “Latest Code Updates” on page 37 for further information.

10.1.1 Backup-Archive Client Code

The backup-archive client is the software piece that goes in the machine that you need to backup. After installing the client code and customizing how it should interact with the ADSM server, you have a working machine ready to send and receive data.

For further assistance on code installation procedures, see the ADSM manual *ADSM V3R1 Installing the Clients* (SH26-4080).

10.1.2 Web Backup-Archive Client

The Web Backup-Archive Client Main Window is shown in Figure 21 on page 151. Its functions are described in the following sections.

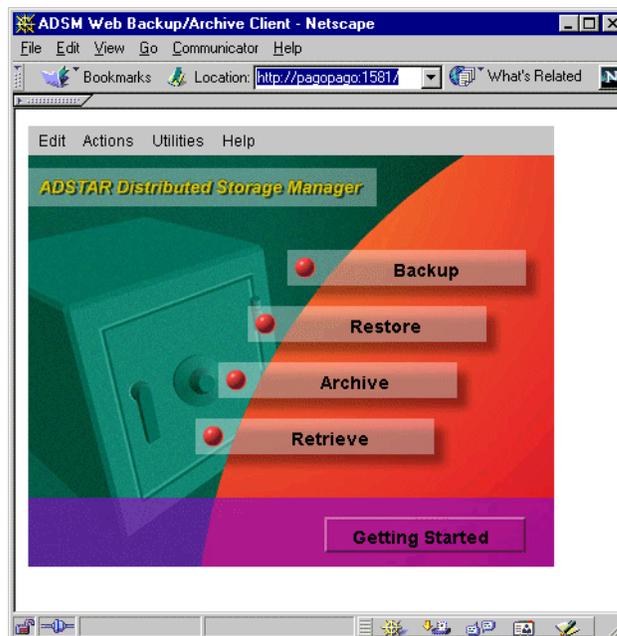


Figure 21. Web Backup-Archive Client Main Window

The Web client consists of new client components that must be installed and configured before the client can be accessed with a Web browser.

The Web client is installed with the backup-archive client package. It is not a separate package as was the Webshell client that it replaces. The Web client consists of two new processes on the client workstation: the client acceptor and the remote client agent.

The client acceptor is an HTTP daemon that serves the Web client Java applet to the Web browser. The name of the executable is DSMCAD. On AIX and other UNIX clients, it should be run as a daemon. For Windows NT it is installed and run as a service. For Windows 95 and 98 it must be executed in a DOS window. For NetWare it is an NLM that should be loaded as part of the NetWare startup.

The remote client agent performs the client functions initiated with the Web client interface. The name of the executable is DSMAGENT. The agent does not have to be running all the time. The acceptor daemon starts the agent when client functions are initiated through the Web client interface.

10.2 Customization

The following describes the different settings for the backup-archive client and the Web backup-archive client.

10.2.1 Backup-Archive Client

This section summarizes the steps you must follow to customize your client installation:

- Setting environment variables
- Defining client options files
- Defining include-exclude lists

For a more detailed explanation see *ADSM V3R1 Installing the Clients* (SH26-4080)

10.2.1.1 Environment Variables

The following are the variables you must update and set in your client environment:

- **PATH:** This is the default search path the operating system uses to locate executable files. Set this to include the fully qualified paths of the ADSM client directories.
- **DSM_CONFIG:** ADSM uses this environment variable to locate the client options file `dsm.opt`. It points to the client user options file for users who create their own personalized options file.
- **DSM_DIR:** ADSM uses this environment variable to locate all other client files. It points to the executable files DSMTCA and DSMSTAT, the resource files, and the `dsm.sys` file (on UNIX only).
- **DSM_LOG:** This points to the directory where you want the `dsmerror.log` file to reside. The error log file contains information about any errors that occur during processing. The error log is intended for IBM service personnel to help you diagnose severe errors.

10.2.1.2 Options File

ADSM includes options that control processing for user sessions. For example, you can use options to inform ADSM which communication method to use, or what format to use for dates.

We provide recommended options files for the various client platforms, as the layout of the supplied options file is not easy to understand. Appendix B.3, “Client Options Files” on page 254 contains our client option files.

The backup-archive client has at least one configuration file, which is divided into the following parts:

- Communication options
- Operational options
- Site-dependent options
- Include-exclude options

Some of the configuration steps that you perform from the client side may need complementary configuration on the ADSM server side.

On all non-UNIX platforms, all options reside in the client options file `dsm.opt` (Preferences file for Macintosh) which resides by default in the ADSM client directory.

On UNIX, you can set options in three different files:

- Client system options file: In the client system options file, a root user sets options that are required to establish communication with an ADSM server, and options that authorize users on your workstation to use ADSM services. A root user can also set options that affect backup and archive processing, and options that affect scheduled services.
- Include-exclude options file: In the include-exclude options file, a root user can set options to exclude specific files from backup services. In addition, a root user can set options to associate specific files with different management classes.
- Client user options file: In the default client user options file, a root user can set options that determine which ADSM server your client node contacts, and that specify the formats to use for date, time, and numbers. A root user can also set options that affect backup, archive, restore, and retrieve processing. In addition, users can also create their own personalized client user options file if they want to use different options. Users can overwrite an option contained in a client user options file by entering a different value for the option with an appropriate ADSM command.

We provide recommended client options files for various client platforms, as the layout of the supplied options file is not easy to understand. Appendix B.3, “Client Options Files” on page 254 contains our client option files. We assume that TCP/IP is the network protocol. Most options can be centrally set using client option sets. We also provide examples in Appendix B.1.2, “Define Client Option Sets” on page 244 for various client platforms.

For more detailed information about client options, refer to the following ADSM manuals:

- *ADSM V3R1 Using the UNIX Backup-Archive Client (SH26-4075)*
- *ADSM V3R1 Using the OS/2 Backup-Archive Client (SH26-4076)*
- *ADSM V3R1 Using the NetWare Backup-Archive Client (SH26-4077)*
- *ADSM V3R1 Using the Windows Backup-Archive Clients (SH26-4078)*

Communication Options

Any ADSM client needs some basic communication settings to send or receive data. You can choose any protocols that are supported in ADSM to use in your environment, provided that both client and server are properly configured for them. Depending on the type of ADSM server, you may need extra help from the system administrator to have all the necessary communication definitions in place. The server is most likely to have many communication protocols configured. If that is the case, you may choose any of those to start communicating with the ADSM server.

The following is a list of communication options you need to set:

- **COMMETHOD:** The COMMETHOD option specifies the communication method you are using to provide connectivity for client-server communication. The ADSM server configuration file (`dsmserv.opt`) must have specified the same communication method with a valid port address, so that the server can accept client requests.
- **TCPSERVERADDRESS:** This option specifies the TCP/IP address for an ADSM server. This is either the IP address or the name of the ADSM server. If you choose to use the name, make sure that there is a name resolution service (DNS, HOSTS, WINS) in place for the TCP/IP configuration in the client machine. Otherwise, you may not be able to connect to the ADSM server. Bear in mind that if you use name resolution to figure out a server address, you will depend on the name resolution protocol to be up and running, which may not be desirable due to network problems that you may experience on the naming resolution services. If you believe that name resolution may not be always available during your backup windows, consider using the IP addresses instead. Although it gives you less flexibility, it will work despite any name resolution service problem.
- **TCPPORT:** The TCPPORT option specifies a TCP/IP port address for an ADSM server. By default, this value is 1500.
- **TCPWINDOWSIZE:** This option specifies the size, in kilobytes, of the TCP/IP sliding window for your client node. This setting is highly operating-system specific. You must only use the allowed values for your TCP/IP implementation.
- **TCPBUFFSIZE:** This option specifies the size, in kilobytes, of the ADSM internal TCP/IP communication buffer. This setting is highly operating-system specific. You must only use the allowed values for your TCP/IP implementation.
- **NODENAME:** This option assigns a new name to your client node if you do not want to use the default. The default is the hostname of your client machine, which we recommend using.

Operational Options

The following list gives some examples of client options for backup, restore, and scheduling services.

- **PASSWORDACCESS:** This option specifies whether you want your ADSM password generated or set as a user prompt. We highly recommend setting `PASSWORDACCESS` to `GENERATE`.
- **REPLACE:** The `REPLACE` option specifies what you want ADSM to do when it restores files that already exist on your workstation. This option applies to the `restore` and `retrieve` commands only.
- **SUBDIR:** The `SUBDIR` option specifies whether you want ADSM to include subdirectories of named directories. This option applies, for example, to `selective`, `restore`, `archive`, and `retrieve`.
- **TAPEPROMPT:** This option specifies whether to wait for a tape to mount if it is required for a backup, archive, restore, or retrieve process, or to prompt you for a choice.
- **SCHEDMODE:** The `SCHEDMODE` option specifies whether you want to use the client-polling mode (your client node periodically asks the ADSM server for scheduled work), or the server-prompted mode (the server contacts your client node when it is time to start a scheduled operation). All communication methods can use the client polling mode, but only TCP/IP can use the server prompted mode.

Site Dependent Options

If your locale is not the USA, then we recommend that you adjust the date, number, and time format options and also the language option in your client options file to support your requirements.

Include-Exclude Options

The include-exclude options may be placed either as part of the client options file `dsm.opt` on non-UNIX platforms, or in a separate file, the include-exclude file on UNIX. On UNIX, you need to add the name of the file in your client system options file `dsm.sys` with the keyword `INCLEXCL` as shown as follows:

```
INCLEXCL /adsm/philip/inclexcl.file
```

The `INCLUDE` option specifies files within a broad group of excluded files that you want to include for backup services. You also use this option to assign a management class either to specific files or to all files to which you have not already assigned a specific management class, and for which you do not want ADSM to use the default management class.

The `EXCLUDE` option excludes files from backup services. When you back up files, any files you exclude are not considered for backup. For example, we recommend you exclude the ADSM installation directory.

Include-exclude options are checked from the bottom up to the top of the list until a match is found. If a match is found, the processing stops and checks whether the option is `INCLUDE` or `EXCLUDE`. If the option is `INCLUDE`, the file is backed up, using the assigned management class. If the option is `EXCLUDE`, the file is not backed up.

See Appendix B.1.2, “Define Client Option Sets” on page 244 for examples of include-exclude lists.

10.2.2 Web Backup-Archive Client

Three client options are important for the Web client. The PASSWORDACCESS option must be set to GENERATE and the password must be generated by running a backup-archive client session. The remote client agent establishes connection to ADSM in the same manner as the backup-archive GUI and command line clients. It requires the generated client password to authenticate a client session when the Web client is used.

The acceptor daemon listens on a TCP/IP port for incoming connections from an administrator's Web browser. By default it listens on port 1581. This default port can be overridden with the new HTTPPORT client option.

The new REVOKEREMOTEACCESS option has two possible values: NONE, which is the default, and ACCESS. If the option is set to NONE, any administrator userid with client access or client owner authority can perform client operations. If the option is set to ACCESS, administrator userids with only client access authority are prevented from performing remote client operations. A pop-up message displayed in the Web browser indicates that the administrator userid being used has insufficient authority. This option does not prevent administrators with client owner or higher authorities from performing client functions. Table 32 on page 156 summarizes the Web client options.

Table 32. Web Client Options

	Option
PASSWORDACCESS	GENERATE
HTTPPORT	1581
REVOKEREMOTEACCESS	NONE

10.3 Interfaces

ADSM client interfaces let you submit and receive information from the ADSM server. In this section we describe how you can start, use, and stop the two available client interfaces: the command line and the GUI. Table 33 on page 156 lists all available backup-archive client interfaces by platform.

Table 33. Backup-Archive Client Interfaces by Platform

Platform	Command Line Interface	GUI	Web backup-archive client
Apple Macintosh		x	
Data General DG/UX	x	x	
Digital UNIX	x	x	x
Hewlett-Packard HP/UX	x	x	x
IBM AIX	x	x	x
IBM OS/2	x	x	x
IBM OS/390 UNIX System Services	x	x V2 only	

Platform	Command Line Interface	GUI	Web backup-archive client
Microsoft Windows 95, 98, and NT/Intel	x	x	x
Microsoft Windows NT DEC Alpha	x	x	
NCR UNIX SVR4	x	x	
NEC EWS-UX/V	x	x	
Novell NetWare	x		x
SCO UNIX	x	x	
Sequent PTX	x	x	
Siemens Nixdorf Reliant UNIX	x	x	
Silicon Graphics IRIX	x	x	x
Sun Solaris	x	x	x

Although you can also have a Web backup-archive client interface, this one is much more related to remote administrative operation of the backup-archive client. You can use the Web interface to remotely execute most of the backup-archive client operations by using an administrative user. For further details on Web backup-archive client interface, see 10.5, “Web Client Usage” on page 179.

There are minor differences between the backup-archive client code among the platforms. For example, in Windows NT you have specific options to handle the NT Registry information, which is not found in any other platform. Despite those specific options, all commands are the same except for the filesystem specification for each platform (in UNIX this is `/usr/lpp`, and in Windows this is `D:\newdir`). Although the examples in this section are based on a UNIX machine, we also give examples of Windows commands.

10.3.1 Command Line

The command line interface has the benefit of being a character mode interface, and thus is well suited for those users who want to type the commands. You may also consider using it when you cannot access the GUI or when you want to automate a backup process by using a batch processing file. You may use it for backup, restore, archive, retrieve operations and to start the ADSM scheduler.

10.3.1.1 Starting a Session

To start the command line interface, you must execute the backup-archive client program by typing `DSMC` in the operating system prompt.

```
kindu:[/usr/lpp/adsm/bin]$ dsmc
ADSTAR Distributed Storage Manager
Command Line Backup Client Interface - Version 3, Release 1, Level 0.6
(C) Copyright IBM Corporation, 1990, 1998, All Rights Reserved.

dsmc>
```

You then get the prompt `dsmc>` to type the backup-archive client commands. When issuing the first command, you may be prompted for your password before the command is executed. This is when you are not using automatic password handling. If you set up your client to use automatic password handling, this is when you set the `PASSWORDACCESS` option to `GENERATE`, the command you issue is executed directly.

Once a command is executed you always return back to the `dsmc>` prompt. This interactive method is also called *loop mode*.

Another method to start an administrative session is the so-called *batch mode*. You simply type `DSMC` together with the actual command that you want to execute. In this mode, the backup-archive client processes the command and returns to the operating system prompt. This is the recommended way to use the `DSMC` command, since it makes automation straightforward. You can have all separate `DSMC` commands in a batch file for automatic processing.

Help is available from the `dsmc>` command prompt by typing `help` and then selecting a number from the list.

```
dsmc> help

The following help topics are available.
Enter the number of the desired help topic or 'q' to quit,
'd' to scroll down, 'u' to scroll up.

    0 - Using Commands
    1 - ARCHIVE
    2 - CANCEL RESTORE
    3 - DELETE ACCESS
    4 - DELETE ARCHIVE
    5 - DELETE FILESPACE
    6 - EVENTBACK
    7 - EVENTREST
    8 - HELP
    9 - INCREMENTAL
   10 - LOOP
   11 - MACRO
   12 - QUERY ACCESS
   13 - QUERY ARCHIVE
   14 - QUERY BACKUP
   15 - QUERY FILESPACE
   16 - QUERY MGMTCLASS
   17 - QUERY RESTORE
   18 - QUERY SCHEDULE

...
```

10.3.1.2 Running Backup Operations

To start backing up a file, directory, or the whole machine, you use the `incremental` command. This command backs up all new or changed files in your default client domain, or in the file systems you specify that are not excluded from backup operations.

Another backup command is the `selective` command. During a selective backup, ADSM sends copies of the files to the server even if they have not changed since the last backup. This might result in having more than one copy of the same file on the server. If that occurs, you might not have as many different file copy versions on the server as you intended. Your version limit might consist of identical files. To avoid that, use the `incremental` command to back up only changed and new files.

Table 34 on page 159 gives you some examples on how to perform a backup operation using a UNIX backup-archive client:

Table 34. Backup Command Examples: UNIX

If you want to perform this then this is the client command
Incremental backup of the client domain	<code>dsmc incremental</code>
Incremental backup of a filesystem (/home)	<code>dsmc incremental /home</code>
Incremental backup of all files in a directory (/home/lucia)	<code>dsmc incremental /home/lucia/</code>
Incremental backup of all files in a directory (/home/lucia) and all its subdirectories	<code>dsmc incremental /home/lucia/ -subdir=yes</code>
Selective backup of a filesystem (/home) and all its subdirectories	<code>dsmc selective /home/ -subdir=yes</code>
Selective backup of all files in a directory (/home/lucia) and all its subdirectories	<code>dsmc selective /home/lucia/ -subdir=yes</code>

Table 35 on page 159 gives you some examples on how to perform a backup operation using a Windows backup-archive client:

Table 35. Backup Command Examples: Windows

If you want to perform this then this is the client command
Incremental backup of the client domain	<code>dsmc incremental</code>
Incremental backup of the C: volume only	<code>dsmc incremental c:</code>
Incremental backup of all files in a directory (C:\work)	<code>dsmc incremental c:\work*.*</code>
Incremental backup of all files in a directory (C:\work) and all its subdirectories	<code>dsmc incremental c:\work*.* -subdir=yes</code>
Selective backup of a directory (C:\work)	<code>dsmc selective c:\work*.*</code>
Selective backup of all files in a directory (C:\work) and all its subdirectories	<code>dsmc selective c:\work*.* -subdir=yes</code>

The following is an example incremental backup operation and its command output:

```
dsmc> incremental /home -subdir=yes

Incremental backup of volume '/home'
Directory-->      512 /home/ [Sent]
Directory-->      512 /home/adsm [Sent]
Directory-->      512 /home/claudia [Sent]
Directory-->      512 /home/daniela [Sent]
Directory-->      512 /home/fred [Sent]
Directory-->      512 /home/guest [Sent]
Directory-->      512 /home/lost+found [Sent]
Directory-->      512 /home/adsm/errors [Sent]
Directory-->      512 /home/adsm/outputs [Sent]
Normal File-->    292 /home/adsm/admin.ksh [Sent]
Normal File-->    384 /home/adsm/console.ksh [Sent]
Normal File-->    1,637 /home/adsm/devconfig.out [Sent]
Normal File-->    59,312 /home/adsm/dsmerror.log [Sent]
Normal File-->      0 /home/adsm/fileexit.out [Sent]
Normal File-->    676 /home/adsm/setup.ksh [Sent]
Normal File-->   19,760 /home/adsm/volhist.out [Sent]
Normal File-->    1,226 /home/adsm/errors/error.out [Sent]
Normal File-->    649 /home/adsm/outputs/hw_definitions [Sent]
Normal File-->    111 /home/adsm/outputs/licenses [Sent]
Normal File-->    181 /home/adsm/outputs/stg_definitions [Sent]
Normal File-->      1 /home/adsm/outputs/various [Sent]
Normal File-->   20,092 /home/claudia/modules.txt [Sent]
Normal File-->   64,761 /home/daniela/adsm.gif [Sent]
Normal File-->      0 /home/fred/xlogfile [Sent]
Successful incremental backup of '/home'

Total number of objects inspected:      24
Total number of objects backed up:      24
Total number of objects updated:         0
Total number of objects rebound:         0
Total number of objects deleted:         0
Total number of objects failed:          0
Total number of bytes transferred:     165.41 KB
Data transfer time:                      0.07 sec
Network data transfer rate:              2,102.91 KB/sec
Aggregate data transfer rate:            110.55 KB/sec
Objects compressed by:                   0%
Elapsed processing time:                  00:00:01
```

10.3.1.3 Running Restore Operations

To restore a file, a directory, or even the whole machine, you need to know two things: *what* you want to restore (file name, directory), and optionally *from when* (point-in-time) if you want to restore a file other than the most recent one.

Important: you do not need to know where the data actually is. When you request a file, ADSM looks into its database and mounts the desired volume. The `restore` command obtains copies of backup versions of your files from an ADSM server.

To restore files, specify the directories or selected files. You also can select the files from a list using the PICK option with the `restore` command, as shown in the following command:

```

dsmc> restore /home/landzett/* -pick

ADSM Scrollable PICK Window - Restore

#      Backup Date/Time      File Size A/I  File
-----
1.    03/10/99  15:02:00      100 A  /home/landzett/.Xauthority
2.    03/10/99  15:02:00     1541 A  /home/landzett/.acrorc
3.    03/10/99  15:02:00      237 A  /home/landzett/.acrosrch
x 4.    03/10/99  15:02:00     2681 A  /home/landzett/.adsmrc
x 5.    03/10/99  15:02:00      512 A  /home/landzett/.dt
x 6.    03/10/99  15:02:00     3968 A  /home/landzett/.dtprofile
x 7.    03/10/99  15:02:00      491 A  /home/landzett/.kshrc
8.    03/10/99  15:02:00       11 A  /home/landzett/.mh_profile
9.    03/10/99  15:02:00      512 A  /home/landzett/.netscape
10.   03/10/99  15:02:00        0 A  /home/landzett/.newsrsrc-news
11.   03/10/99  15:02:00      327 A  /home/landzett/.profile
12.   03/10/99  15:02:00     2494 A  /home/landzett/.sh_history
13.   03/10/99  15:02:00      512 A  /home/landzett/Mail
14.   03/10/99  15:02:00     1988 A  /home/landzett/admin.log

0-----10-----20-----30-----40-----50-----60-----7
<U>=Up <D>=Down <T>=Top <B>=Bottom <R#>=Right <L#>=Left
<C#>=Goto Line # <#>=Toggle Entry <+>=Select All <->=Deselect All
<#:#+>=Select A Range <#:#->=Deselect A Range <O>=Ok <C>=Cancel
pick>

```

You can restore files to the directory from which you backed them up, or you can restore files to a different location. Depending on what you want to restore, you may need to use the PRESERVEPATH option, which specifies how much of the source path you want to preserve and append to the destination path specified with the restore operation.

A file copy can be in one of three states: active, inactive, or expired. An active file copy is the most current copy of the file, an inactive file copy is a previous copy of the file, and an expired file copy is a copy to be removed from the ADSM server. Only active versions are considered for restore unless you use the INACTIVE or LATEST options: The INACTIVE option instructs ADSM to restore an inactive backup if an active one is not available, and the LATEST option restores the most recent backup version of a file, even if the backup is inactive.

Table 36 on page 161 gives you some examples on how to perform a restore operation using a UNIX backup-archive client:

Table 36. Restore Command Examples: UNIX

If you want to perform this then this is the client command
Restore a single file (/home/myfile1)	dsmc restore /home/myfile1
Restore a directory and all files from that level (/home)	dsmc restore "/home/"
Restore all files in a directory (/home) with all its subdirectories	dsmc restore "/home/" -subdir=yes
Restore a full directory (/home to /temp) with all subdirectories and write full path on destination	dsmc restore /home/myfile1 /temp -preservepath=complete -subdir=yes

Table 37 on page 162 gives you some examples on how to perform a restore operation using a Windows backup-archive client:

Table 37. Restore Command Examples: Windows

If you want to perform this then this is the client command
Restore a single file (C:\work\myfile1)	<code>dsmc restore c:\work\myfile1</code>
Restore a directory and all files from that level (C:\work)	<code>dsmc restore c:\work\</code>
Restore all files from a directory (C:\work) with all its subdirectories	<code>dsmc restore c:\work\ -subdir=yes</code>
Restore a full subdirectory (C:\work to D:\temp) with all subdirectories, and write full path on destination	<code>dsmc restore c:\work\ d:\temp -preservepath=complete -subdir=yes</code>

The following is an example restore operation and its command output. The restored data does not exist on the client workstation prior to the restore operation.

```
dsmc> restore /home/ -subdir=yes

Restore function invoked.

ANS1247I Waiting for files from the server...
Restoring          512 /home/ [Done]
Restoring          512 /home/adsm [Done]
Restoring          512 /home/claudia [Done]
Restoring          512 /home/daniela [Done]
Restoring          512 /home/fred [Done]
Restoring          512 /home/guest [Done]
Restoring          512 /home/lost+found [Done]
Restoring          512 /home/adsm/errors [Done]
Restoring          512 /home/adsm/outputs [Done]
Restoring          292 /home/adsm/admin.ksh [Done]
Restoring          384 /home/adsm/console.ksh [Done]
Restoring         1,637 /home/adsm/devconfig.out [Done]
Restoring         59,312 /home/adsm/dsmerror.log [Done]
Restoring           0 /home/adsm/fileexit.out [Done]
Restoring          676 /home/adsm/setup.ksh [Done]
Restoring        19,760 /home/adsm/volhist.out [Done]
Restoring          1,226 /home/adsm/errors/error.out [Done]
Restoring          649 /home/adsm/outputs/hw_definitions [Done]
Restoring          111 /home/adsm/outputs/licenses [Done]
Restoring          181 /home/adsm/outputs/stg_definitions [Done]
Restoring           1 /home/adsm/outputs/varioues [Done]
Restoring        20,092 /home/claudia/modules.txt [Done]
Restoring        64,761 /home/daniela/adsm.gif [Done]
Restoring           0 /home/fred/xlogfile [Done]

Restore processing finished.

Total number of objects restored:      24
Total number of objects failed:       0
Total number of bytes transferred:    165.49 KB
Data transfer time:                   0.35 sec
Network data transfer rate:           472.06 KB/sec
Aggregate data transfer rate:         83.60 KB/sec
Elapsed processing time:               00:00:01
```

10.3.1.4 Running Archive Operations

The `archive` command archives a single file, selected files, or all files in a directory and its subdirectories on an ADSM server. You use this command to save files that you want to keep in their present condition for a long period of time, and to control them by expiration date only. There is no version controlling for archives. You can also use archives to release storage space on your workstation and delete files as you archive them using the `DELETEFILES` option. Typically, you may use archives to save information that has either a legal requirement (for example, account information, annual reports, billing information, or annual customer reports); or an internal audit requirement (for example, application logs, user activity information, or employee files).

When you use the `archive` command with the `ARCHMC` option, you can select an available management class in the active policy set that has an archive copygroup. This enables you to select the retention period (according to the retention period in days specified in the archive copygroup) for all the data you are archiving. You can also use the `DESCRIPTION` option to group data together, so that they can be later retrieved without having to locate the files. This feature is called packaging. The description is case-sensitive. The default includes the date and time stamp of the archive operation.

Table 38 on page 163 gives you some examples on how to perform an archive operation using a UNIX backup-archive client:

Table 38. Archive Command Examples: UNIX

If you want to perform this then this is the client command
Archive all files from a directory (/home)	<code>dsmc archive /home/ -subdir=yes</code>
Archive files from different directories (/home and /temp)	<code>dsmc archive /home/ /temp/ -subdir=yes</code>
Archive files from a directory (/home) to a management class called Y1 (which is one year retention)	<code>dsmc archive /home/ -archmc=y1</code>
Create a new archive package called "DANCLA-FILES" from a directory (/home)	<code>dsmc archive /home/ -subdir=yes -description="DANCLA-FILES"</code>
Add files from a directory (/temp) to an archive package called "DANCLA-FILES"	<code>dsmc archive /temp/ -subdir=yes -description="DANCLA-FILES"</code>

Table 39 on page 163 gives you some examples of how to perform an archive operation using a Windows backup-archive client:

Table 39. Archive Command Examples: Windows

If you want to perform this then this is the client command
Archive all files from a directory (C:\work)	<code>dsmc archive c:\work\ -subdir=yes</code>
Archive files from different directories (C:\work and C:\temp)	<code>dsmc archive c:\work\ c:\temp\ -subdir=yes</code>
Archive files from a directory (C:\work) to a management class called Y1 (which is one year retention)	<code>dsmc archive c:\work\ -archmc=y1</code>
Create a new package called "DANCLA-FILES" from a directory (C:\work)	<code>dsmc archive c:\work\ -subdir=yes -description="DANCLA-FILES"</code>

If you want to perform this then this is the client command
Add files from a directory (C:\temp) to an archive package called "DANCLA-FILES"	<code>dsmc archive c:\temp\ -subdir=yes -description="DANCLA-FILES"</code>

The following is an example archive operation and its command output:

```

dsmc> archive /home/ -subdir=yes -description="1998files" -archmc=y1

Archive function invoked.

Directory-->          512 /home/ [Sent]
Directory-->          512 /home/adsm [Sent]
Directory-->          512 /home/claudia [Sent]
Directory-->          512 /home/daniela [Sent]
Directory-->          512 /home/fred [Sent]
Directory-->          512 /home/guest [Sent]
Directory-->          512 /home/lost+found [Sent]
Directory-->          512 /home/adsm/errors [Sent]
Directory-->          512 /home/adsm/outputs [Sent]
Normal File-->        292 /home/adsm/admin.ksh [Sent]
Normal File-->        384 /home/adsm/console.ksh [Sent]
Normal File-->        1,637 /home/adsm/devconfig.out [Sent]
Normal File-->        59,312 /home/adsm/dsmerror.log [Sent]
Normal File-->         0 /home/adsm/fileexit.out [Sent]
Normal File-->        676 /home/adsm/setup.ksh [Sent]
Normal File-->       19,760 /home/adsm/volhist.out [Sent]
Normal File-->        1,226 /home/adsm/errors/error.out [Sent]
Normal File-->        649 /home/adsm/outputs/hw_definitions [Sent]
Normal File-->        111 /home/adsm/outputs/licenses [Sent]
Normal File-->        181 /home/adsm/outputs/stg_definitions [Sent]
Normal File-->         1 /home/adsm/outputs/various [Sent]
Normal File-->       20,092 /home/claudia/modules.txt [Sent]
Normal File-->       64,761 /home/daniela/adsm.gif [Sent]
Normal File-->         0 /home/fred/xlogfile [Sent]
Archive processing of '/home/' finished without failure.

Total number of objects inspected:      24
Total number of objects archived:      24
Total number of objects updated:        0
Total number of objects rebound:       0
Total number of objects deleted:        0
Total number of objects failed:         0
Total number of bytes transferred:     165.41 KB
Data transfer time:                     0.07 sec
Network data transfer rate:             2,329.70 KB/sec
Aggregate data transfer rate:           137.81 KB/sec
Objects compressed by:                  0%
Elapsed processing time:                 00:00:01

```

10.3.1.5 Running Retrieve Operations

The `retrieve` command obtains copies of archived files from the ADSM server. You can specify either selected files or whole directories to retrieve archived files. Use options such as `description` that allow you to search for the descriptions assigned to the files when they were archived.

Table 40 on page 165 gives you some examples on how to perform an retrieve operation using a UNIX backup-archive client:

Table 40. Retrieve Command Examples: UNIX

If you want to perform this then this is the client command
Retrieve a single file (/home/myfile1)	<code>dsmc retrieve /home/myfile1</code>
Retrieve files from a directory (/home)	<code>dsmc retrieve /home</code>
Retrieve files from a directory (/home) with a description "project1998" and save them into another location (/temp)	<code>dsmc retrieve /home/ /temp/ -subdir=yes -description="project1998"</code>

Table 41 on page 165 gives you some examples on how to perform a retrieve operation using a Windows backup-archive client:

Table 41. Retrieve Command Examples: Windows

If you want to perform this then this is the client command
Retrieve a single file (C:\work\myfile1)	<code>dsmc retrieve c:\work\myfile1</code>
Retrieve files from a directory (C:\work)	<code>dsmc retrieve c:\work\ -subdir=yes</code>
Retrieve files from a directory (C:\work) with a description "project1998" and save them into another location (C:\temp)	<code>dsmc retrieve c:\work\ c:\temp\ -subdir=yes -description="project1998"</code>

Note:

1. The description field is case-sensitive.
2. The same restore rules applies for the retrieve operations, when using the PRESERVEPATH option.

The following is an example retrieve operation and its command output. The retrieved data does not exist on the client workstation prior to the retrieve operation.

)

```
dsmc> retrieve /home/ -subdir=yes -description="1998files"

Retrieve function invoked.

Retrieving          512 /home/adsm [Done]
Retrieving          512 /home/claudia [Done]
Retrieving          512 /home/daniela [Done]
Retrieving          512 /home/fred [Done]
Retrieving          512 /home/guest [Done]
Retrieving          512 /home/lost+found [Done]
Retrieving          292 /home/adsm/admin.ksh [Done]
Retrieving          384 /home/adsm/console.ksh [Done]
Retrieving         1,637 /home/adsm/devconfig.out [Done]
Retrieving        59,312 /home/adsm/dsmerror.log [Done]
Retrieving          512 /home/adsm/errors [Done]
Retrieving          512 /home/ [Done]
Retrieving           0 /home/adsm/fileexit.out [Done]
Retrieving           0 /home/fred/xlogfile [Done]
Retrieving          512 /home/adsm/outputs [Done]
Retrieving          676 /home/adsm/setup.ksh [Done]
Retrieving        19,760 /home/adsm/volhist.out [Done]
Retrieving         1,226 /home/adsm/errors/error.out [Done]
Retrieving          649 /home/adsm/outputs/hw_definitions [Done]
Retrieving          111 /home/adsm/outputs/licenses [Done]
Retrieving          181 /home/adsm/outputs/stg_definitions [Done]
Retrieving           1 /home/adsm/outputs/various [Done]
Retrieving        20,092 /home/claudia/modules.txt [Done]
Retrieving        64,761 /home/daniela/adsm.gif [Done]

Retrieve processing finished.

Total number of objects retrieved:      24
Total number of objects failed:         0
Total number of bytes transferred:    165.49 KB
Data transfer time:                    0.19 sec
Network data transfer rate:            861.27 KB/sec
Aggregate data transfer rate:          107.72 KB/sec
Elapsed processing time:                00:00:01
```

10.3.1.6 Stopping a Session

If the ADSM backup-archive client is running in loop mode and resting at the `dsmc>` prompt, then you have to enter `quit` to end the session. This terminates the connection with the ADSM server and returns to the calling program, which is normally the operating system prompt.

```
dsmc> quit
kindu:[/usr/lpp/adsm/bin]$
```

Note:

If you want to terminate a running backup-archive session, you can cancel the operation by pressing CTRL-C in the session window or ask the ADSM administrator to issue the `cancel session` command. If you choose to cancel it by using CTRL-C while there is a restore operation in progress, the ADSM administrator needs to cancel the restartable restore by issuing the `cancel restore` command from the ADSM server console or from an administrative client.

10.3.2 GUI

The GUI is a user-friendly interface with unique features for the end-user. You may use it for all ADSM operations, but not for the ADSM scheduler.

10.3.2.1 Starting a Session

To start the GUI, you must execute the backup-archive client program by typing DSM in the operating system prompt. The GUI has a slightly different appearance on Windows and UNIX machines. Despite the appearance, UNIX and Windows backup-archive clients have the same functionality. Although we will use the Windows GUI in the next examples to show how to run client operations, keep in mind that the same functions apply for the UNIX version as well.

Figure 22 on page 167 shows the main screen of the backup-archive client GUI interface on Windows:

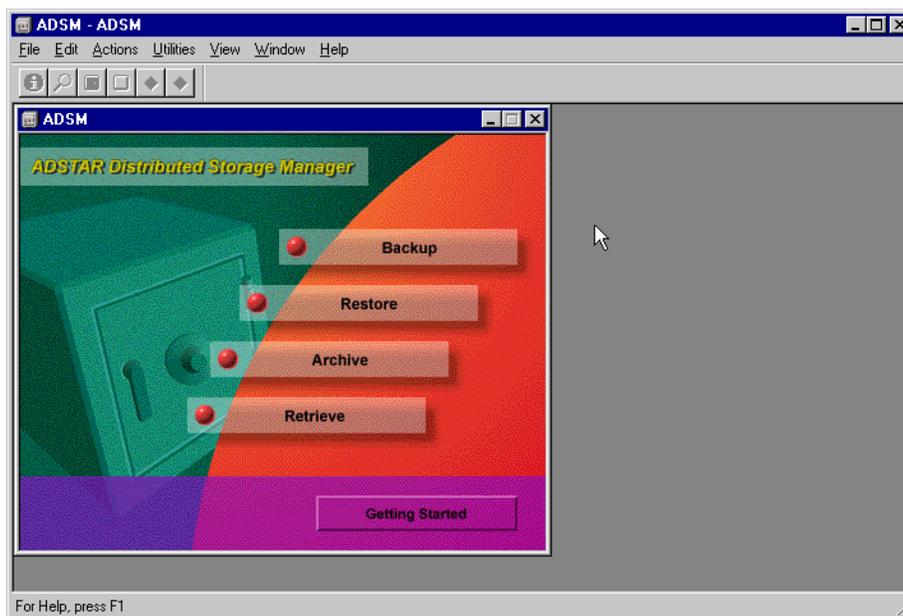


Figure 22. Backup-Archive Client GUI: Windows

Figure 23 on page 168 shows the main screen of the backup-archive client GUI interface on UNIX:



Figure 23. Backup-Archive Client GUI: UNIX

Depending on the features of your UNIX machine, for example, if it has the Common Desktop Environment (CDE) installed, you can also add an option in the workplace screen so that it is easier to start ADSM.

10.3.2.2 Running Backup Operations

To open the Backup window, click on **Backup** in the main window as shown in Figure 22 on page 167. Figure 24 on page 168 shows an example of a Backup window with some files selected.

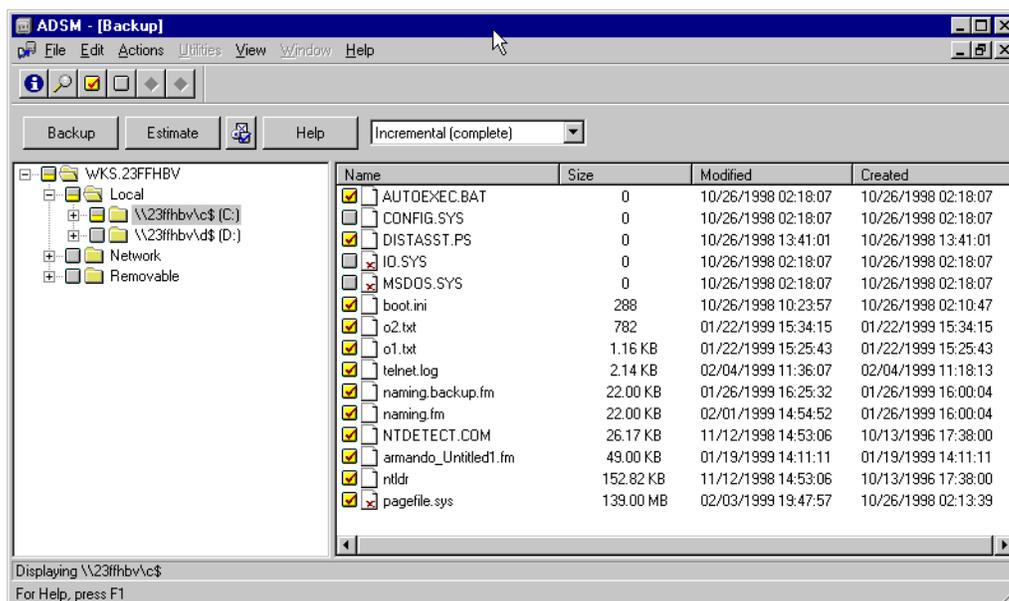


Figure 24. GUI Backup Window

To back up some of your files:

1. In the Backup window, expand the directory tree. Click on the plus sign (+) next to a directory or drive.
2. Click the selection boxes next to the drives or directories you want to back up.
3. Be sure that the drop-down list near the top of the Backup window displays the Incremental (Complete) choice.
4. Click on **Backup**. The Backup Status window appears.
5. You can view the progress of the backup in the Backup Status window. When it is finished, you get a completion dialog box.

10.3.2.3 Running Restore Operations

To open the Restore window, click on **Restore** in the main window as shown in Figure 22 on page 167. Figure 25 on page 169 shows an example of a Restore window with some files selected.

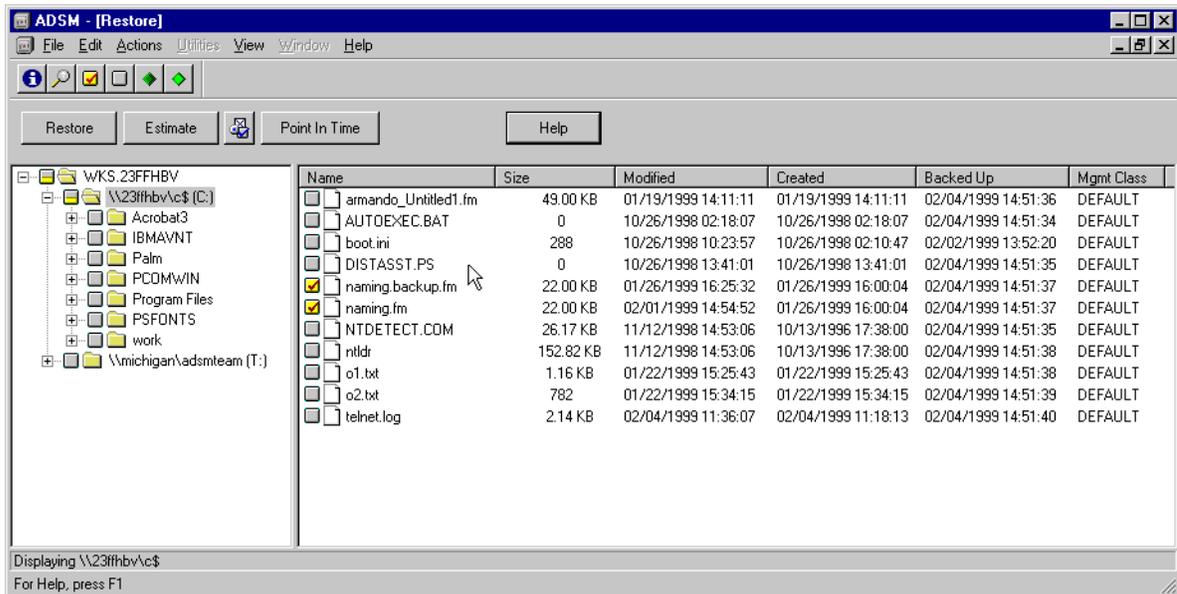


Figure 25. GUI Restore Window

The Restore window contains objects which represent directories and files that you previously backed up to the ADSM server. Use this window to select the objects that you want to restore. You can also request ADSM to run an estimate for the restore or change processing options.

There are several restore operations that you can perform:

- Restore backup versions from directory tree
- Restore backup versions by name
- Restore backup versions by filtering the directory tree
- Restore someone else's backup versions
- Restore inactive backup versions
- Run a point-in-time restore

- Work with restartable restore sessions

Restore Backup Versions from Directory Tree

1. In the Restore window, expand the directory tree. Click the plus sign (+) next to a directory or drive.
2. Click the selection boxes next to the files or directories that you want to restore.
3. Click on **Restore**. The Restore Destination window appears.
4. Fill in the Restore Destination window, then click on **Restore** to start the restore.
5. You can view the progress of the restore in the Restore Status window. When it is finished, you get a completion dialog box.

Restore Backup Versions by Name

1. In the Restore window, expand the directory tree. Click the plus sign (+) next to a directory or drive.
2. Highlight the drive or directory you want to work with.
3. Click the Search tool on the tool bar.
4. Enter your search criteria in the Find Files window, then click on **Search**.
5. When a list of files appears in the Matching Files window, click the selection boxes next to the files you want to restore. Your selections appear in the Restore by Tree window.
6. Click on **Restore**. The Restore Destination window appears.
7. Fill in the Restore Destination window, then click on **Restore** to start the restore.
8. You can view the progress of the restore in the Restore Status window.

Restore Backup Versions by Filtering the Directory Tree

1. In the Restore window, expand the directory tree. Click the plus sign (+) next to a directory or drive.
2. Highlight the drive or directory you want to work with.
3. Click the Search tool on the tool bar.
4. Enter your criteria in the Find Files window, then click the Filter button. Close the window. The directory tree displays only those directories that contain the files that match your criteria.
5. Click the selection boxes next to the files or filtered directories that you want to restore.
6. Click on **Restore**. The Restore Destination window appears.
7. Fill in the Restore Destination window, then click on **Restore** to start the restore.
8. You can view the progress of the restore in the Restore Status window.

Restore Someone Else's Backup Versions

1. From the main window, access the other user's stored data
 1. Select the Access Another User choice from the Utilities menu. The Access Another User window appears.

2. Type the user's node name in the Node Name field.
 3. Click the Set button. You can display another user's stored data only if the other user has authorized you to do so in the Authorization List window.
2. Restore the other user's backup versions using one of the following methods:
 - Restoring backup versions by name
 - Restoring backup versions from the directory tree
 - Restoring backup versions by filtering the directory tree

Restore Inactive Backup Versions

1. In the Restore window, select the Display Active/Inactive Files choice from the View menu. The inactive file icon appears next to file names of inactive backup versions in the file list.
2. Expand the directory tree. Click the plus sign (+) next to a directory or drive.
3. Click the selection boxes next to the directories or files that you want to restore.
4. Click the Restore button. The Restore Destination window appears.
5. Fill in the Restore Destination window, then click on **Restore** to start the restore.
6. You can view the progress of the restore in the Restore Status window.

Run a Point-in-Time Restore

1. In the Restore window, click on **Point in Time Restore**.
2. Fill in the Point in Time Restore window and click on **OK**. The point in time that you specified appears in the Point in Time display field in the Restore window.
3. Display the objects you want to restore. To display the objects, you can search for an object by name, filter the directory tree, or work with the directories in the directory tree.
4. Click the selection boxes next to the objects you want to restore.
5. Click on **Restore**. The Restore Destination window appears.
6. Fill in the Restore Destination window, then click on **Restore** to start the restore.
7. You can view the progress of the restore in the Restore Status window.

Work with Restartable Restore Sessions

To restart a restore session that was interrupted:

1. From the main window, select the Restartable Restores choice from the Actions menu. If there are any restartable restore sessions available, the Restartable Restores window appears.
2. Select a restore session that you want to restart from the list.
3. Click on **Restart**.

To delete a restore session that was interrupted:

1. From the main window, select the Restartable Restores choice from the Actions menu. If there are any restartable restore sessions available, the Restartable Restores window appears.

2. Select an object that you want to delete from the list.
3. Click on **Delete**.

10.3.2.4 Running Archive Operations

To open the Archive window, click on **Archive** in the main window as shown in Figure 22 on page 167. Figure 26 on page 172 shows an example of an Archive window with some files selected.

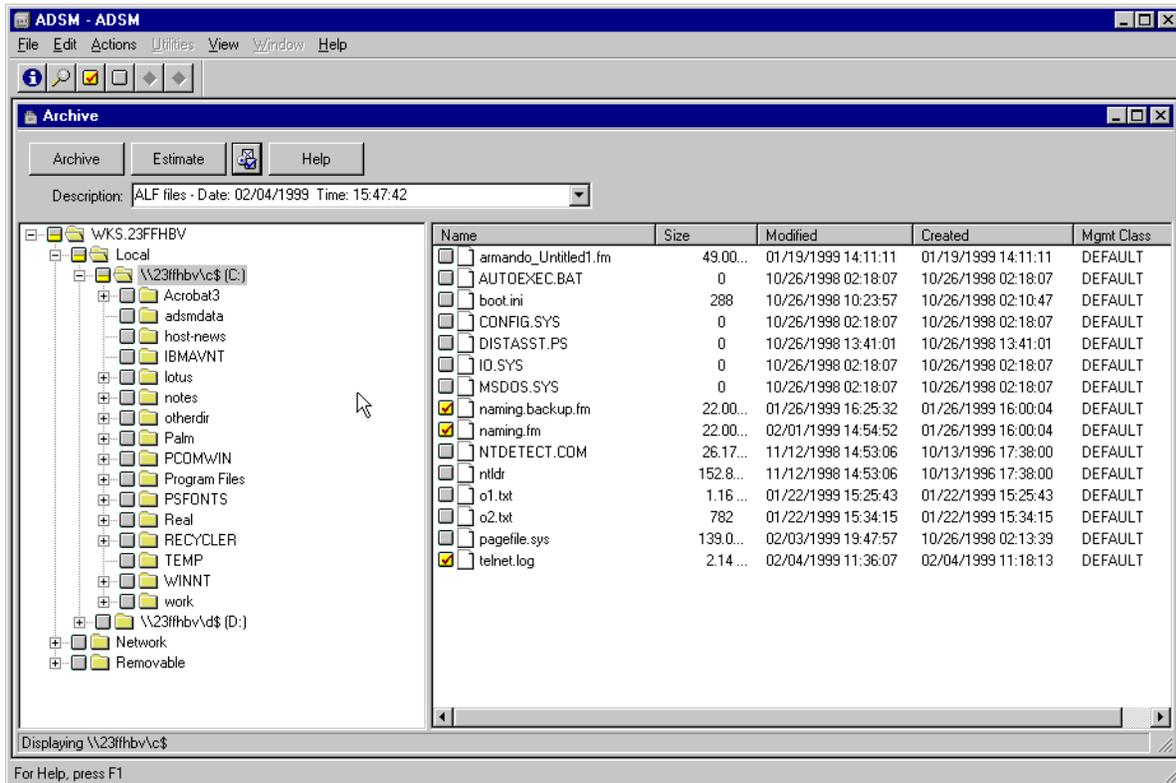


Figure 26. GUI Archive Window

The Archive window contains objects which represent directories and files on your workstation. Use this window to select the directories and files you want to archive.

ADSM requires that you assign an archive description for all archived files. An archive description identifies data through a meaningful description that you can use later to identify files and directories. You can enter as many as 255 characters to describe your archived data. If a description is not entered, ADSM assigns a default archive description.

When you select the archive function from the backup-archive GUI, ADSM displays a list of all previously used archive descriptions. You can use these displayed archive descriptions on future archives.

You can archive specific files or entire directories from the directory tree. You can assign a unique description for each group of files you archive. This group of files is called an archive package. Archiving files and directories in such a way makes

retrieval easier. You can archive and retrieve files or directories as one entity, retrieve individual files, add files to an existing package using the archive description, or delete files from a package.

To archive your files:

1. In the Archive window, expand the directory tree. Click the plus sign (+) next to a directory or drive.
2. Click the selection boxes next to the files or directories that you want to archive.
3. Type a description, accept the default description, or select an existing description for your archive package in the Description combination box.
4. Click on **Archive**. The Archive Status window appears as shown in Figure 26 on page 172.
5. You can view the progress of the archive in the Archive Status window. When it is finished, you get a completion dialog box.

10.3.2.5 Running Retrieve Operations

To open the Retrieve window, click on **Retrieve** in the main window as shown in Figure 22 on page 167. Figure 27 on page 173 shows an example of a Retrieve window with some files selected

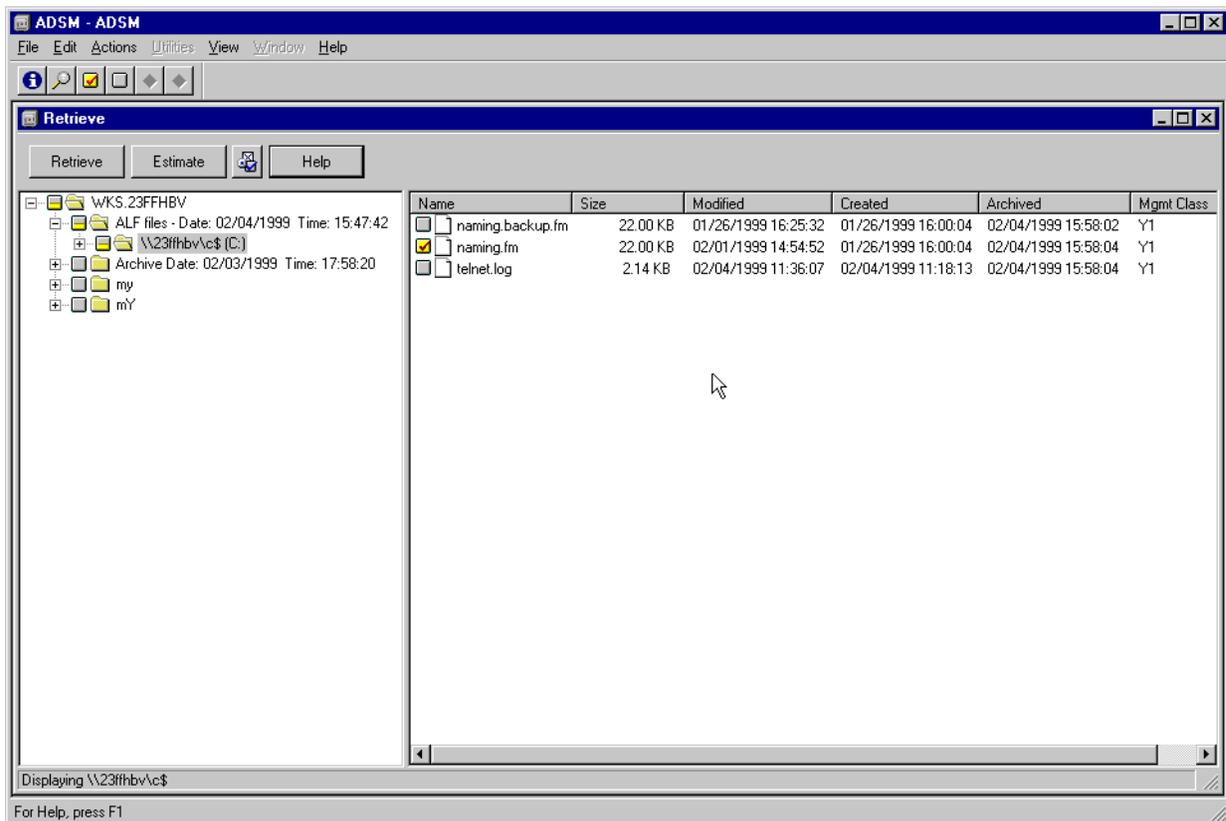


Figure 27. GUI Retrieve Window

The Retrieve window contains objects which represent objects that you archived to the ADSM server. Use this window to select the objects you want to retrieve.

To retrieve your archive copies:

1. In the Retrieve window, expand the directory tree. Click the plus sign (+) next to an object you want to expand. If you are connected to the ADSM Version 3 server, the objects on the tree are grouped by archive package description.
2. Click the selection boxes next to the objects that you want to retrieve in the tree view or file view.
3. Click on **Retrieve**. The Retrieve Destination window appears.
4. Fill in the Retrieve Destination window, then click on **Retrieve** to start the operation.
5. You can view the progress of the retrieve in the Retrieve Status window. When it is finished, you get a completion dialog box.

You can also retrieve another user's archive copies to your local drives if the other user has authorized you to do so.

To retrieve another user's archive copies:

1. From the main window, access the other user's stored data
 1. Select the Access Another User choice from the Utilities menu. The Access Another User window appears.
 2. Type the user's node name in the Node Name field.
 3. Click on **Set..**
2. Retrieve the other user's archive copies using the previous method explained.

10.3.2.6 Stopping a Session

To stop a GUI session, you must select File -> Exit from the pull-down menu in the main window.

10.4 Client Scheduler

An ADSM administrator can schedule ADSM to perform tasks automatically on a regular basis. For example, you can automatically back up files at the end of each day, or archive some of the files every Friday. This procedure, known as central scheduling, is a cooperative effort between the server and the client node. You associate clients with one or more schedules that are part of the policy domain maintained in the ADSM database. You use the central scheduling on the server, and you start the client scheduler on the workstation. Once you start the client scheduler, further intervention is not necessary.

Note:

The schedule start time is based on the server clock, not the client clock.

With client scheduling, you can also:

- Display information about available schedules.
- Display information about work that the schedule has completed.
- Modify scheduling options in the client options file.

Note:

Install the ADSM command-line client and ensure that the communication software is running before you start the client scheduler. You cannot start the client scheduler from the ADSM GUI.

10.4.1 Starting the Client Scheduler

You can start the ADSM client scheduler either manually or with each system start. ADSM does not recognize changes to the client options files while the client scheduler is running. If you make changes to these files while the client scheduler is running, and you want ADSM to use the new values immediately, stop the client scheduler and restart it.

UNIX Users:

The client scheduler can only be started by an ADSM-authorized user. This is a user who has administrative authority for the ADSM client on a workstation. This user changes passwords, performs open registrations, and deletes filespace.

10.4.1.1 Manual Start

You may perform a manual startup for the ADSM scheduler when you accidentally stopped it, or even when you want to force it to run immediately.

To start the client scheduler on your client node and connect to the server scheduler, enter the `dsmc schedule` command from the operating system prompt. You cannot start the scheduler in loop mode.

```

kindu:[usr/lpp/adsm/bin]$> dsmc schedule
ADSTAR Distributed Storage Manager
Command Line Backup Client Interface - Version 3, Release 1, Level 0.1
(C) Copyright IBM Corporation, 1990, 1997, All Rights Reserved.

Node Name: AIXCLIENT1
Session established with server AIX1: AIX-RS/6000
  Server Version 3, Release 1, Level 2.1
  Server date/time: 01/25/1999 10:52:44  Last access: 01/25/1999 10:50:36

Querying server for next scheduled event.
Next operation scheduled:
-----
Schedule Name:      INCR_AIXCLIENT1
Action:             Incremental
Objects:
Options:
Server Window Start: 23:35:00 on 07/30/1999
-----
Schedule will be refreshed in 12 hours.

Time remaining until execution: 11:59:56

```

When you start the client scheduler, it runs continuously until you close the window, thus end the process, or log off your system.

UNIX

To run the schedule command in background and to keep the client scheduler running even if you log off your UNIX system, enter the following command from the UNIX shell:

```

kindu:[usr/lpp/adsm/bin]$> nohup dsmc schedule 2> /dev/null &

```

Windows

You may create a batch file, for example, `sched.bat`, to perform the client scheduler start, so that there will be no need to type the ADSM command. The following commands should accomplish this:

```

rem
rem BATCH file for ADSM client scheduler start
rem
cd c:\Program Files\IBM\adsm\baclient
set DSM_CONFIG=c:\Program Files\IBM\adsm\baclient\dsm.opt
set DSM_DIR=c:\Program Files\IBM\adsm\baclient
dsmc schedule -password=secret

```

Figure 28. Batch File for Scheduler Start on Windows

You can also specify the `dsmc schedule` command in your startup folder so that every time you start Windows, the client scheduler starts.

10.4.1.2 Automatic Start

There are many cases where it is desirable that the scheduler is automatically started. This may be the case if you have an enterprise policy to have ADSM machines backed up at some point in the day.

UNIX

If you want to start the client scheduler when you start your operating system, set the PASSWORDACCESS option parameter to GENERATE in your client system options file, and include the following command in the /etc/inittab file:

```
adsm::once:/usr/bin/dsmc sched > /dev/null 2>&1 # ADSM scheduler
```

The scheduler starts the next time the system is rebooted.

For OS/390 UNIX system services, add the following entry to the /etc/rc file:

```
_BPX_JOBNAME='ADSM' /usr/adsm/dsmc schedule &
```

Windows NT

You can register the ADSM schedule as a service under Windows NT using the DSMCUTIL program, so that it becomes part of the operating system startup procedure.

The following is an example of how to install the client scheduler named "ADSM Scheduler" for node CLADAN on an Windows NT machine. For further information on the syntax of this utility program check the `dsmcutil.txt` file in the ADSM backup-archive client directory:

```

C:\Program Files\IBM\adsm\baclient>dsmcutil install /name:"ADSM Scheduler"
/optfile:"c:\program files\ibm\adsm\baclient\dsm.opt"
/clientdir:"c:\program files\ibm\adsm\baclient" /node:CLADAN
/password:secret /eventlog:yes /autostart:yes

ADSM Windows NT Client Service Configuration Utility
Command Line Interface Version 2.00.b
Last Updated Nov 18 1998

Command: Install ADSM Client Service
Machine: CLADAN (Local Machine)

Installing ADSM Client Service:

      Machine           : CLADAN
      Service Name      : ADSM_SCHED
      Client Directory  : c:\program files\ibm\adsm\baclient
      Automatic Start   : Yes
      Logon Account     : LocalSystem

The service was successfully installed.

Creating Registry Keys ...

Updated registry value 'ImagePath' .
Updated registry value 'EventMessageFile' .
Updated registry value 'TypesSupported' .
Updated registry value 'OptionsFile' .
Updated registry value 'EventLogging' .
Updated registry value 'ClientNodeName' .
Updated registry value 'ADSMClientKey' .

Generating registry password ...
Authenticating password with ADSM for node CLADAN....

Connecting to ADSM Server via client options file 'c:\program
files\ibm\adsm\baclient\dsm.opt' ...

Password authentication successful.

The Registry password for node CLADAN has been updated .

Starting the 'ADSM_SCHED' service .....

The service was successfully started.

```

10.4.2 Stopping the Scheduler

If you want to terminate a running client scheduler, you can cancel the operation by pressing CTRL-C in the session window, or you can ask the ADSM administrator to issue the `cancel session` command. If you choose to cancel it by using CTRL-C while there is a restore operation in progress, the ADSM administrator needs to cancel the restartable restore by issuing the `cancel restore` command from the ADSM server console or from an administrative client.

On Windows NT, you can terminate the client scheduler by stopping its service. To change the status of the scheduler service, you have to go to Settings -> Control Panel -> Services, select the ADSM client scheduler service and stop it.

10.5 Web Client Usage

The ADSM Version 3 Web backup-archive client is a Java client interface that enables authorized users to perform remote operations on a backup-archive client system.

10.5.1 Remote Client GUI

The Web client is a Java applet that provides a remote client GUI for ADSM Version 3 backup-archive clients. Authorized users can access the Web client remotely using a Web browser. The graphic is the initial client hub window displayed in the Web browser when the Web client is accessed.

The Web client can be used from any workstation running a Web browser with support for Java 1.1.5 or higher. Netscape 4.06 has the required Java support. Netscape 4.03, 4.04, and 4.05 requires a JDK 1.1 upgrade for this support. This upgrade can be obtained from the Netscape homepage. Microsoft's Internet Explorer 4.01 has the required Java support. If the browser does not have the correct level of Java support, the client hub window is displayed as a grey box.

The Web client is supported only on ADSM Version 3 servers.

10.5.2 Client Functions

The Web client supports most of the Version 3 backup-archive client functions. It can be used to perform backup, restore, archive, and retrieve operations. For restore operations, both active and inactive files can be restored, and the point-in-time restore function is supported. The functions not currently supported include cross client restores, filespace deletion, archive file deletion, and the file find and search functions.

The Windows NT Web client can back up and restore the NT registry. The NetWare Web client can backup and restore the NetWare 3.x bindery and NetWare 4.x Directory Services (NDS) and, for the first time, provides a NetWare client GUI.

The Web client is another client interface. It can be used only to perform backup-archive client operations. If used from a Web browser on a remote workstation, no local access to backup or archive data is provided on that remote workstation. Data cannot be restored or retrieved locally; it can only be restored to or retrieved to the client workstation that owns the data. A Web browser can also be used locally on the client workstation to invoke the Web client as an alternative interface to the backup-archive client command line interface or GUI.

10.5.3 Access Authorization

The Web client is a new interface in addition to the client's existing GUI and command line interface. Use of the Web client interface is authenticated whenever backup, restore, archive, or retrieve Web client functions are performed. Authentication of the Web client interface is separate, and independent, from authentication between the client node and the server.

An administrator userid is required to use the Web client. This administrator userid, and associated password, is used to authenticate that the user has sufficient authority to perform remote client functions. Two new administrative authorities are provided to enable this authentication:

- Client owner
- Client access

These new authorities are available with the Version 3.1.2.1 PTF level of the ADSM server and can be used to enable usage of the Web client interface for backup-archive client owners and helpdesk personnel.

The Web backup-archive client connects to an ADSM client via the Web to present an interface to remotely perform the functions of the backup-archive client on that client. Using a Web browser means that you actually do not have to have any ADSM code on the Web browser workstation.

This interface is very useful because it provides the basic backup, restore, archive, and retrieve functions remotely and without authorization problems.

10.5.4 Starting the Web Client

The Web client consists of two processes on the client workstation: the client acceptor and remote client agent.

10.5.4.1 Client Acceptor

The client acceptor is an HTTP daemon that serves the Web client Java applet to the Web browser. The name of the executable is DSMCAD.

On AIX and other UNIX clients, it should be run as a daemon.

To have the Web client automatically available when the system is started, you have to edit the /etc/inittab file and add the following entry in the /etc/inittab:

```
dsmcad:2:once:/usr/lpp/adsm/bin/dsmcad >/dev/console 2>&1
```

For Windows NT it is installed and run as a service. By default, the installation procedure registers the Web Client Acceptor as a manual service. Therefore, if you want to make it an automatic service, you have to go to Settings -> Control Panel -> Services, and change the startup setting of the Web Client Acceptor to automatic.

For Windows 95 and 98, the DSMCAD program must be executed in a DOS window.

For NetWare, DSMCAD is an NLM that should be loaded as part of the NetWare startup.

10.5.4.2 Remote Client Agent

The remote client agent performs the client functions initiated with the Web client interface. The name of the executable is DSMAGENT. The agent does not have to be running all the time. The acceptor daemon starts the agent when client functions are initiated through the Web client interface.

10.5.4.3 Web Client

To access the Web backup-archive client, start your Web browser. In the Location or Address field, enter the URL:

http://<machinename>:1581

where <machinename> is the TCP/IP DNS name or dotted IP address of your ADSM client machine and 1581 is the TCP/IP port number set by the HTTPPORT option.

Figure 29 on page 181 shows you the main Web client screen:

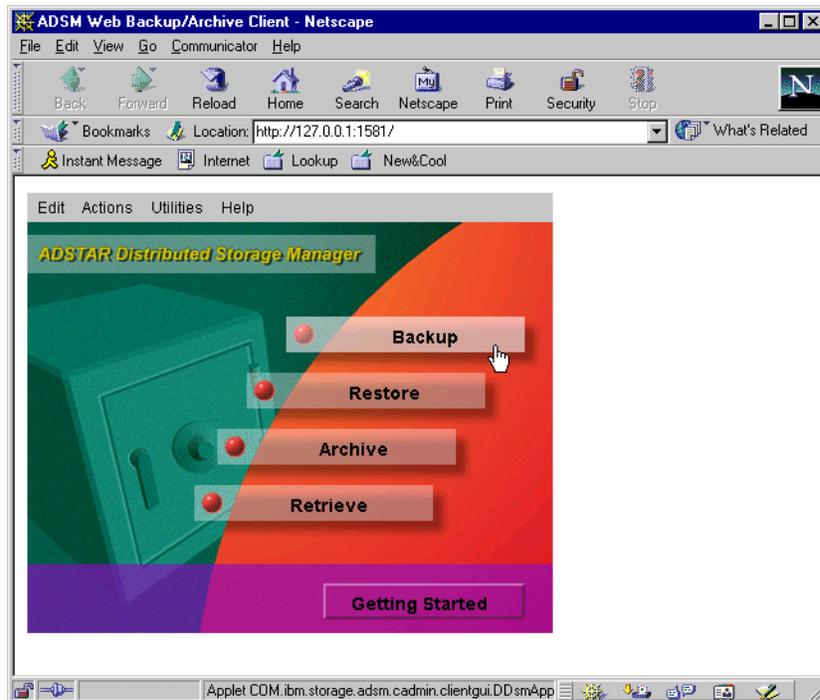


Figure 29. Web Backup-Archive Client Main Window

When you try to connect to a client by using the Web client interface, ADSM validates the administrative ID you supply. You are then granted access to the Web client interface and the client is connected to the ADSM server for backup, archive, restore, and retrieve operations. Figure 30 on page 181 shows a login dialog box for a Web client session:



Figure 30. Web Client Login Dialog Box

10.5.5 Stopping the Remote Agent Services

In general, there is no need to stop the remote agent services manually.

On Windows NT, if you wish to disable administrators from accessing the Web client, you can prevent the service from starting by disabling it. To do that, you have to go to Settings -> Control Panel -> Services and then change the startup settings to disabled.

In a UNIX environment, you can stop the DSMCAD process by using the `KILL` command:

```
kindu:[/work/adsm]$ ps -ef | grep dsmcad
root 12506 13446  0 02:28:06 pts/2  0:00 grep dsmcad
root 24078      1  0 09:46:46      -  0:00 ./dsmcad

kindu:[/work/adsm]$ kill -15 24078

kindu:[/work/adsm]$ ps -ef | grep dsmcad
root 14302 13446  1 02:28:27 pts/2  0:00 grep dsmcad
```

Chapter 11. Scheduling

ADSM scheduling is divided into two different categories: administrative and client scheduling. The two categories differ in two key areas:

- Execution location: An administrative schedule performs some action on the ADSM server, while the client schedule can only execute on the ADSM client.
- Domain privilege: An administrative schedule can only be managed by an administrator with system privilege, while the client schedule can be managed by an administrator with policy privileges in only one domain. The granularity provided by this feature can be very useful when distributing management control across a large enterprise.

For both types of schedules, there are four key pieces of information:

- A command that needs to be executed
- When the command executes
- How long the command takes to complete
- How often the command needs to be repeated

The command that you run may be an incremental backup (client schedule) or a storage pool migration (administrative schedule) which should be run every day at a particular time. You also have to estimate how long the command will run so that you can synchronize your schedules and balance load on the server.

This chapter provides scheduling rules of thumb and recommendations that will give you a good base configuration.

11.1 The Wheel of Life

Figure 31 on page 184 shows the recommended series of operations that should occur in an ADSM environment and the sequence in which those operations should occur. The circle represents a clock and the tick marks indicate the hours of the day. The actual start time and duration of the various operations are subject to your scheduling requirements, but not the sequence.

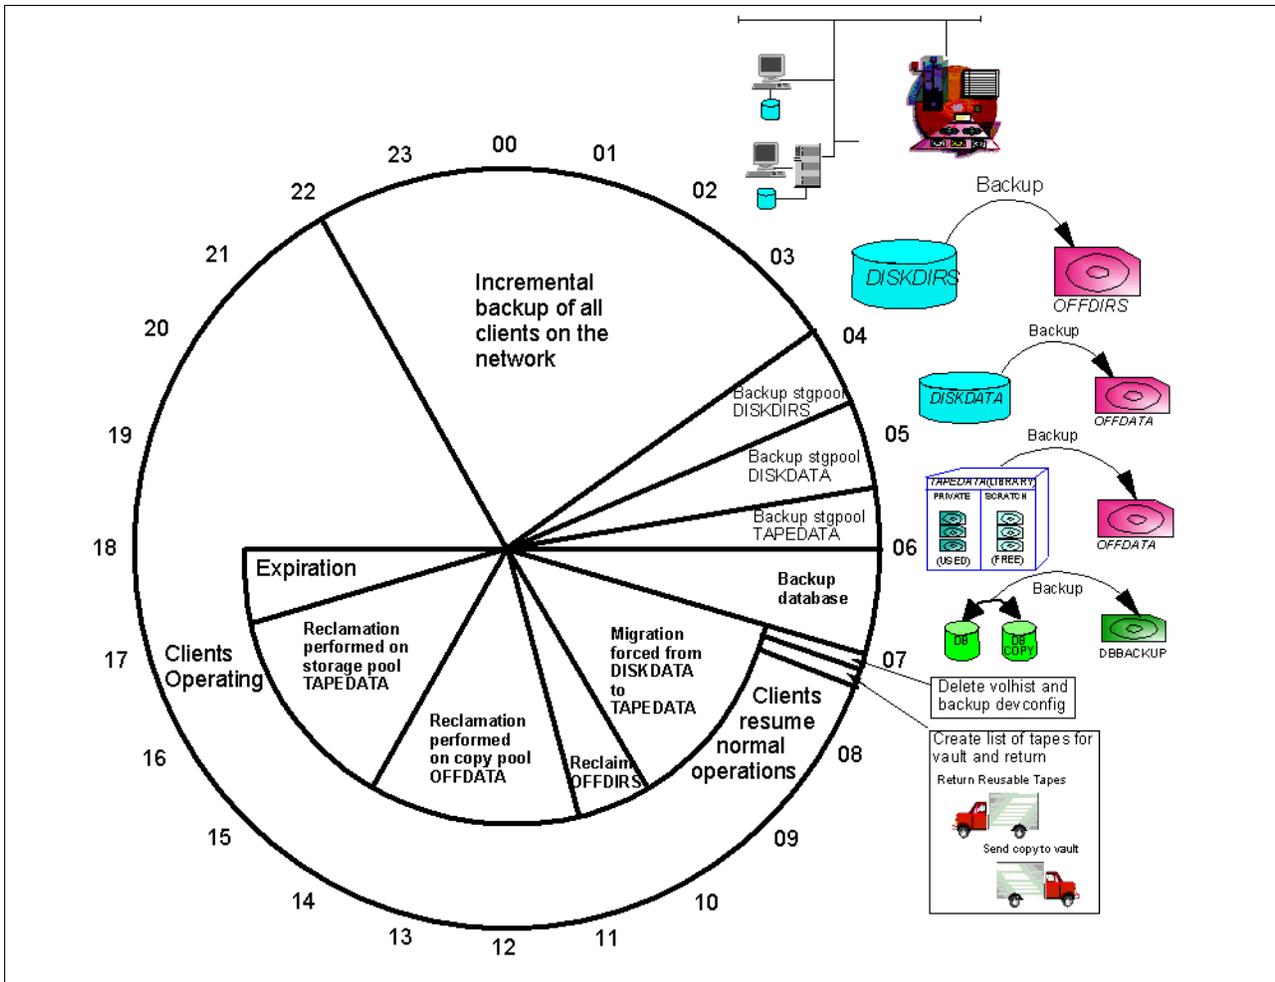


Figure 31. Scheduling of Operations

11.2 Administrative Schedules

An administrative schedule is a directive to trigger some sort of action on the ADSM server. It consists of a command or sequence of commands and some details on when the actions should happen. Any actions that are used on a regular basis to manage the ADSM environment should be defined as an administrative schedule.

In 1.2, “Redbook Support Material” on page 2, we show how to load a predefined macro into ADSM. The macro `mac.schedules`, which we provide to define schedules for our redbook environment, is shown in Appendix B.1.4, “Define Schedules” on page 246.

The following command illustrates the syntax for an administrative schedule. This command defines a schedule named `BACKUP_DATABASE` that specifies a full database backup, and that will start at 7 a.m. every morning.

```

adsm> define schedule backup_database type=administrative \
cont> cmd="backup db deviceclass=offsite type=full" starttime=07:00 active=yes \
cont> description="Database backup"
ANR2577I Schedule BACKUP_DATABASE defined.

```

We have assembled a series of administrative schedules that should help you set up an ADSM environment that will minimize user intervention while providing a high level of data availability.

The first step in setting up your administrative schedules is knowing which commands to run. Table 42 on page 185 shows our recommendations for commands that should be scheduled and where to find more information on them. The table also includes our estimates on various other factors based on a typical implementation. Client schedules are not considered administrative, but they are included to emphasize that all of the other scheduling in your environment should hinge on your client schedules. This is a recommended set of schedules with all of the specifics required to execute in your environment if you follow the recommendations made in this redbook.

Table 42. Recommended Schedules

#	Function	Task	Duration (Hrs)	Task Dependency	Refer to:
1	Backup Client Data to ADSM Server	Schedule Client Backup Windows	6.0	Site Requirements	11.3
2	Copy Backup Data to Offsite Storage Media	Backup Storage Pool DISKDIRS	0.5	Task 1 Complete	11.2.1
3		Backup Storage Pool DISKDATA	1.0	Task 1 Complete	
4		Backup Storage Pool TAPEDATA	1.0	Task 1 Complete	
5		Backup Database	0.5	Task 2,3,4 Complete	
6	Volume History File Management	Delete Volume History Data for Database Backup Volumes	0.25	Task 5 Complete	11.2.2
7	Prepare Disk Storage Pools for Next Backup Window	Migrate data from DISKDATA to TAPEDATA	3.0	Task 5 Complete	11.2.3
8	Defragment Tape Volumes in Sequential Storage Pools	Reclaim volumes from OFFDIRS	1.0	None	11.2.4
9		Reclaim volumes from OFFDATA	3.0	None	
10		Reclaim volumes from TAPEDATA	3.0	None	
11	Database Space Management	Expire Inventory	1.0	None	11.2.5

#	Function	Task	Duration (Hrs)	Task Dependency	Refer to:
12	Enable Statistics	Audit Licenses	0.25	None	11.2.6

11.2.1 Defining an Offsite Backup Schedule

When the client backups are complete, all of the data that has changed in the primary storage pools should be copied to a copy storage pool. The copy storage pool tapes and the database backup tapes should be taken to another site, called the offsite location. These tapes are referred to collectively as the offsite backups, and will be used to restore your environment in the event of a catastrophic failure or disaster.

11.2.1.1 Flow of Events

The order of execution of these tasks should be:

1. Check that there are no running sessions with the clients. When you back up the primary storage pools and database, you want to make sure that you are capturing an up to date copy. If there is still a client backing up its data, you should wait until it is done before starting the offsite backups. Use an SQL command to check for running sessions.
2. Back up the primary storage pools to their copy storage pool. See 5.9.3, "Back Up Storage Pool" on page 95 for more information.
 1. Back up the DISKDIRS storage pool to the OFFDIRS storage pool
 2. Back up the DISKDATA storage pool to the OFFDATA storage pool
 3. Back up the TAPEDATA storage pool to the OFFDATA storage pool
3. Back up the database. See 4.8, "Database Backup" on page 62 for more information.

Since all of these commands are going to take a variable amount of time to execute, and they should be run in sequence, it makes sense to use a server script to run them. This will guarantee that each command will complete its tasks before another command is started.

11.2.1.2 Defining the Server Script

Figure 32 on page 186 displays the commands that we would type into an administrative command line interface to accomplish the offsite backup tasks. We can use a server script to run these commands and check return codes to determine if we can proceed with the remaining commands.

```
select * from sessions where upper(session_type)='NODE'
backup stgpool diskdirs offdirs wait=yes
backup stgpool diskdata offdata wait=yes
backup stgpool tapedata offdata wait=yes
backup db devclass=offsite type=full scratch=yes
```

Figure 32. Administrative Commands to Back Up Storage Pools and Database

Figure 33 on page 187 shows the finished script for our recommended configuration. In 1.2, “Redbook Support Material” on page 2, we show how to load a predefined macro into ADSM. The macro `mac.scripts`, which we provide to define a script to backup storage pools and the database to an offsite location, is shown in Appendix B.1.5, “Define Server Scripts” on page 247.

```

/* Script file for automation of the offsite storage tasks */

/* Make sure that any previous copies of the script are removed */
del script offsite_backups

/* Backup all primary storage pools followed by the database */
def script offsite_backups desc='Backup all primary storage pools &
database'
upd script offsite_backups /* -----*/
upd script offsite_backups /* Script Name: OFFSITE_BACKUPS */
upd script offsite_backups /* Description: Backup all primary storage */
upd script offsite_backups /* pools followed by the */
upd script offsite_backups /* database. If there are */
upd script offsite_backups /* active node sessions, don't*/
upd script offsite_backups /* run now, but reschedule to */
upd script offsite_backups /* check and run again in 20 */
upd script offsite_backups /* minutes. */
upd script offsite_backups /* The command will backup */
upd script offsite_backups /* three storage pools called */
upd script offsite_backups /* diskdirs, diskdata and */
upd script offsite_backups /* tapdata to backup copypools */
upd script offsite_backups /* called offdata and offdirs. */
upd script offsite_backups /* Example: run offsite_backups */
upd script offsite_backups /* -----*/
upd script offsite_backups ' select * from sessions where -'
upd script offsite_backups " upper(session_type)='NODE' "
upd script offsite_backups ' if (rc_ok) goto reschedule'
upd script offsite_backups ' ba stg diskdirs offdirs wait=yes'
upd script offsite_backups ' ba stg diskdata offdata wait=yes'
upd script offsite_backups ' ba stg tapedata offdata wait=yes'
upd script offsite_backups ' ba db dev=offsite type=full scratch=yes'
upd script offsite_backups ' exit'
upd script offsite_backups ' reschedule:'
upd script offsite_backups ' del sch retry_offsite_backups type=a'
upd script offsite_backups ' def sch retry_offsite_backups t=a cmd="run
offsite_backups -'
upd script offsite_backups ' " active=yes startt=NOW+0:20 peru=o'

```

Figure 33. Script File for Offsite Backups

The name of the script is completely arbitrary, but it should be representative of the function performed by the script. You should check that the script name you choose is not already used by someone else. This can be accomplished by using the following command:

```

adsm> select * from script_names where NAME='REDBOOK_OFFSITE'
ANR2034E SELECT: No match found using this criteria.
ANS8001I Return code 11.

```

Once the script is loaded into the server, you need to define a schedule that will run this script for you every day at 4 a.m. The macro `mac.schedules`, which we provide to define recommended schedules in our redbook environment, is shown in Appendix B.1.4, “Define Schedules” on page 246 and includes the following schedule definition command:

```

adsm> define schedule redbook_offsite type=administrative \
cont> cmd="run REDBOOK_OFFSITE" starttime=04:00 \
cont> description="Backup primary storage pools & database"
ANR2577I Schedule REDBOOK_OFFSITE defined.

adsm> query schedule redbook* type=administrative

```

*	Schedule Name	Start Date/Time	Duration	Period	Day
-	REDBOOK_OFFSITE	03/18/1999 04:00:00	1 H	1 D	Any

11.2.2 Defining the Volume History Schedule

Every volume that is used by ADSM, including the volume identifier for the database backups, is tracked in the volume history database. The database volume information is important because it tells you which volume holds your most recent database backup. In the event of a disaster, this information will be very valuable. The volume history information is periodically copied out to a volume history file that you can specify with the `VOLUMEHISTORY` option in the `dsmserv.opt` file. We recommend that you have two copies of the volume history file, in case one becomes unusable.

There are two recommended schedules that affect the volume history file: `DELETE_VOLHIST` and `BACKUP_VOLHIST`.

11.2.2.1 DELETE_VOLHIST

When you back up the database, the previous database backups become obsolete and should be returned to scratch volume status for reuse. The way to accomplish this is to delete the volume history record of this volume using the `delete volhistory` command. We recommend that you delete any database backups that are older than five days, since we back up the database every day. To define a delete volume history schedule, we define a schedule called `DELETE_VOLHIST` specifying a type of `TYPE=DBBACKUP`. The backslash (`\`) indicates to ADSM to continue the command on the next line. For example:

```

adsm> define schedule delete_volhist type=administrative \
cont> cmd="delete volhistory type=dbbackup todate=today-5" starttime=07:00 \
cont> description="Delete volume history information for database backups"
ANR2577I Schedule DELETE_VOLHIST defined.

adsm> query schedule delete* type=administrative

```

*	Schedule Name	Start Date/Time	Duration	Period	Day
-	DELETE_VOLHIST	01/28/99 07:00:00	1 H	1 D	Any

11.2.2.2 BACKUP_VOLHIST

Some operations, like the `delete volhistory` command, will update the volume history database, but will not trigger the server to update the volume history files. To ensure that the volume history data is always current, you should schedule a daily backup of the volume history database using the `backup volhistory` command. Use these commands to schedule a daily backup of the volume history database:

```
adsm> define schedule backup_volhist type=administrative \  
cont> cmd="backup volhistory" starttime=07:05 \  
cont> description="Backup volume history file"  
  
ANR2577I Schedule BACKUP_VOLHIST defined.  
  
adsm> query schedule backup_volh* type=administrative
```

*	Schedule Name	Start Date/Time	Duration	Period	Day
-	BACKUP_VOLHIST	01/28/99 07:05:00	1 H	1 D	Any

11.2.3 Defining a Migration Schedule

Migration of your primary storage pools during client backup can slow down the client backup sessions and impact the scheduling of other server processes. We recommend that you avoid migration during client backup by scheduling primary storage pool migration outside of the client backup window. This will only prevent migration during client backup if the storage pools are correctly sized to hold the daily backup data from your clients.

To schedule a migration of a storage pool, you change the storage pool high and low migration thresholds to zero using the `update stgpool` command. This will force the storage pools to migrate immediately to their next level of storage pool. The migration process will continue until the storage pool is empty or until you reset the migration thresholds to a non-zero value. It is important not to leave the migration thresholds at zero after they are finished migrating. If more data is sent to these storage pools while they are set to zero, the data will be migrated immediately, defeating the purpose of having the storage pool.

Two schedules are required, one to start the migration and one to stop it. To define the migration schedules to start and stop storage pool migration, do the following:

```
adsm> define schedule migration_start type=administrative \  
cont> cmd="update stgpool diskdata hi=0 lo=0" starttime=07:00 \  
cont> description="Start migration on DISKDATA storage pool"  
ANR2577I Schedule MIGRATION_START defined.  
  
adsm> define schedule migration_stop type=administrative \  
cont> cmd="update stgpool diskdata hi=70 lo=30" starttime=10:00 \  
cont> desc="Stop migration on DISKDATA storage pool"  
ANR2577I Schedule MIGRATION_STOP defined.  
  
adsm> query schedule migration* type=administrative
```

*	Schedule Name	Start Date/Time	Duration	Period	Day
-	MIGRATION_START	01/28/99 07:00:00	1 H	1 D	Any
-	MIGRATION_STOP	01/28/99 10:00:00	1 H	1 D	Any

11.2.4 Defining a Reclamation Schedule

Tapes in the sequential access storage pools will eventually get fragmented due to file expiration and node removal. This can cause many tapes to have very little data on them. A reclamation process will take all of this data and put it on a few tapes thus returning many tapes to scratch status for reuse.

To start a reclamation process on a sequential storage pool, use the `update stgpool` command to set the RECLAIM parameter for the storage pool to a value between 50 and 99. This value represents the percentage of reclaimable space on the volume which is the opposite of the volume utilization. That is to say, if the output from a `query volume` command shows a volume as 20% utilized, then its reclaimable space is 80%. Reclamation stops when the process is complete or the reclaim parameter is set to 100.

Each sequential storage pool should have a schedule to start and stop reclamation. We have three sequential storage pools in our recommended configuration so we need six schedules to start and stop the reclamations. To set up the reclamation schedules, do the following:

```
adsm> define schedule RECLAIM_OFFDIRS_START type=administrative \  
cont> cmd="update stgpool offdirs reclaim=75" starttime=10:00 \  
cont> description="Start reclaim on the OFFDIRS storage pool"  
ANR2577I Schedule RECLAIM_OFFDIRS_START defined.  
  
adsm> define schedule RECLAIM_OFFDIRS_STOP type=administrative \  
cont> cmd="update stgpool offdirs reclaim=100" starttime=11:00 \  
cont> description="Stop reclaim on the OFFDIRS storage pool"  
ANR2577I Schedule RECLAIM_OFFDIRS_STOP defined.  
  
adsm> define schedule RECLAIM_OFFDATA_START type=administrative \  
cont> cmd="update stgpool offdata reclaim=75" starttime=11:00 \  
cont> description="Start reclaim on the OFFDATA storage pool"  
ANR2577I Schedule RECLAIM_OFFDATA_START defined.  
  
adsm> define schedule RECLAIM_OFFDATA_STOP type=administrative \  
cont> cmd="update stgpool offdata reclaim=100" starttime=14:00 \  
cont> description="Stop reclaim on the OFFDATA storage pool"  
ANR2577I Schedule RECLAIM_OFFDATA_STOP defined.  
  
adsm> define schedule RECLAIM_TAPEDATA_START type=administrative \  
cont> cmd="update stgpool tapedata reclaim=75" starttime=14:00 \  
cont> description="Start reclaim on the TAPEDATA storage pool"  
ANR2577I Schedule RECLAIM_TAPEDATA_START defined.  
  
adsm> define schedule RECLAIM_TAPEDATA_STOP type=administrative \  
cont> cmd="update stgpool tapedata reclaim=100" starttime=17:00 \  
cont> description="Stop reclaim on the TAPEDATA storage pool"  
ANR2577I Schedule RECLAIM_TAPEDATA_STOP defined.  
  
adsm> query schedule reclaim* type=administrative
```

*	Schedule Name	Start Date/Time	Duration	Period	Day
-	RECLAIM_OFFDATA- _START	01/28/99 11:00:00	1 H	1 D	Any
	RECLAIM_OFFDATA- _STOP	01/28/99 14:00:00	1 H	1 D	Any
	RECLAIM_OFFDIRS- _START	01/28/99 10:00:00	1 H	1 D	Any
	RECLAIM_OFFDIRS- _STOP	01/28/99 11:00:00	1 H	1 D	Any
	RECLAIM_TAPEDAT- A_START	01/28/99 14:00:00	1 H	1 D	Any
	RECLAIM_TAPEDAT- A_STOP	01/28/99 17:00:00	1 H	1 D	Any

11.2.5 Defining an Expiration Schedule

Managing the amount of space used by your database will be very important as your environment grows. The client backup process handles half of the task by marking expired data references while backing up client data. The second half of the task is handled by the `expire inventory` command. This command will remove the marked entries from the database.

The ADSM server configuration file defaults to a 24 hour cycle for expiration. This would result in daily expiration processing beginning 24 hours from when you start the server. This command can be computationally intensive and it is best scheduled during a quiet server time to minimize its impact on server operations. Thus, we recommend scheduling the `expire inventory` command daily:

```
adsm> define schedule expire_inventory type=administrative \  
cont> cmd="expire inventory" startt=17:00 \  
cont> description="Inventory expiration"  
ANR2577I Schedule EXPIRE_INVENTORY defined.  
  
adsm> query schedule expire* type=administrative
```

*	Schedule Name	Start Date/Time	Duration	Period	Day
-	EXPIRE_INVENTORY	01/28/99 17:00:00	1 H	1 D	Any

11.2.6 Defining a Licensing Audit Schedule

It is helpful to be able to check statistics for the ADSM server usage by using the `query auditoccupancy` command. This command returns some valuable information on client data usage in the ADSM environment.

```
adsm> query auditoccupancy  
License information as of last audit on 02/17/1999 at 15:15:29.
```

Node Name	Backup Storage Used (MB)	Archive Storage Used (MB)	Space-Managed Storage Used (MB)	Total Storage Used (MB)
CLIENT	0	0	0	0
ANDY	182	41	0	223
PAGOPAGO	10	89	0	99
PHIL	0	0	0	0
TORONTO	0	0	0	0
NORFOLK	0	0	0	0
COCOS	0	0	0	0
ABC	0	0	0	0
PUTNEY	0	0	0	0
RYDE	0	0	0	0

The data returned from this command is updated when the `audit licenses` command is run on the ADSM server. We recommend running this command daily using the administrative schedule below:

```
adsm> define schedule audit_license type=administrative \  
cont> cmd="audit licenses" starttime=00:00 \  
cont> description="Audit licenses"  
ANR2577I Schedule AUDIT_LICENSE defined.  
  
adsm> query schedule audit* type=administrative
```

*	Schedule Name	Start Date/Time	Duration	Period	Day
-	AUDIT_LICENSE	01/28/99 00:00:00	1 H	1 D	Any

11.3 Client Schedules

A client schedule is a directive to trigger an action on a group of ADSM client machines. It is different from an administrative schedule in that it specifies that an action be performed on the ADSM client. The client scheduling system consists of a server portion and a client portion. The server part is integrated into the ADSM process and is responsible for defining the schedule parameters and associating nodes with the schedule. The client scheduler is a separate process on the ADSM client and provides communication between the server and client. See 10.4, “Client Scheduler” on page 174 for further information on the client portion of the client scheduling system.

An ADSM client machine with an active scheduler component can be scheduled to perform any of the following actions:

- Backup (incremental/selective)
- Archive
- Restore
- Retrieve
- Operating system command
- Macro

The most common use of client schedules is to implement an automatic incremental backup process for a group of machines.

At this point in the configuration of the ADSM system, you will have installed the ADSM client code on your clients and defined them to a policy domain. You will also have activated the scheduler service on the client so that it is ready to perform the actions that we schedule for it. The remaining activities are concentrated on the server and its scheduling facility. To schedule an action on the client you will need to do the following:

1. Define a client schedule
2. Associate a client with a schedule
3. Verify the client schedules

11.3.1 Defining a Client Schedule

In 1.2, “Redbook Support Material” on page 2, we show how to load a predefined macro into ADSM. The macro `mac.schedules`, which we provide to define schedules for our redbook environment, is shown in Appendix B.1.4, “Define Schedules” on page 246.

The `define schedule` command defines a schedule to the ADSM server. It does not contain any reference to a client; that operation will be handled by the `define association` command. To define a client schedule, you will need the following information:

- Policy domain name: The schedule will be created within this policy domain and can only be associated with nodes within the same domain. This will allow you to manage which administrators have control of these schedules by limiting access to the domain that owns them.
- Schedule name: This can be any name you choose, up to 30 characters.

- Start time: The time of day that you want the schedule to trigger your action.
- Duration: The scheduler can only initiate the schedule within the time period specified with the duration parameters.

The following command illustrates the syntax for a client schedule. This command defines a schedule named BACKUP_NIGHTLY in the ABCCORP policy domain. This schedule will trigger an incremental backup (the default action for a client schedule) at 10 p.m. every night. If there is a problem contacting the client, the scheduler will keep trying until 1a.m.

```
adsm> define schedule abccorp backup_nightly starttime=22:00 duration=3 \
cont> durunits=hours description="Nightly backup schedule for SERVER domain"
ANR2500I Schedule BACKUP_NIGHTLY defined in policy domain SERVER.
```

Client schedules are very dependent on site requirements, but we have some recommendations that are applicable for a wide range of environments.

Our recommended policy configuration includes two policy domains, so we must define a client backup schedule for each. The parameters for each schedule are very similar except for the name which is chosen to reflect the domain. The schedules start at 10 p.m. and will continue attempting to start incremental backups until 1 a.m. The actual commands and the output look like this:

```
adsm> define schedule server server_nightly starttime=22:00 duration=3 \
cont> durunits=hours description="Nightly backup schedule for SERVER domain"
ANR2500I Schedule SERVER_NIGHTLY defined in policy domain SERVER.

adsm> define schedule workstn workstn_nightly starttime=22:00 duration=3 \
cont> durunits=hours description="Nightly backup schedule for WORKSTN domain"
ANR2500I Schedule WORKSTN_NIGHTLY defined in policy domain WORKSTN.

adsm> query schedule
```

Domain	* Schedule Name	Action	Start Date/Time	Duration	Period	Day
SERVER	SERVER_NIGHTLY	Inc Bk	01/30/99 22:00:00	3 H	1 D	Any
WORKSTN	WORKSTN_NIGHTLY	Inc Bk	01/30/99 22:00:00	3 H	1 D	Any

11.3.2 Associating a Client with a Schedule

Once the schedule is defined, we need to specify which nodes belong to it using the `define association` command. We must associate nodes to a schedule within their own domain. The nodes must be registered to the server before running this command, but they do not necessarily have to be in contact with the server. The actual commands and the output look like this:

```

adsm> define association server server_nightly putney
ANR2510I Node PUTNEY associated with schedule SERVER_NIGHTLY in policy domain
SERVER.

adsm> define association workstn workstn_nightly cocos
ANR2510I Node COCOS associated with schedule WORKSTN_NIGHTLY in policy
domain WORKSTIN.

adsm> query association

Policy Domain Name: SERVER
Schedule Name: SERVER_NIGHTLY
Associated Nodes: PUTNEY

Policy Domain Name: WORKSTIN
Schedule Name: WORKSTN_NIGHTLY
Associated Nodes: COCOS

```

11.3.3 Verifying the Client Schedules

The schedules can be checked out by using the query event command. This will show you if the nodes and schedules are all set up correctly. The actual command and its output look like this:

```

adsm> query event * *

```

Scheduled Start	Actual Start	Schedule Name	Node Name	Status
03/17/1999 22:00:00		SERVER_NIGHTLY	PUTNEY	Future
03/17/1999 22:00:00		WORKSTN_NIGHTLY	COCOS	Future

Chapter 12. Operations

ADSM, as any other software system, has to be maintained and operated. This is to check that the tasks which are implemented are still working. To accomplish this, you may need to consider how many people will be dedicated to ADSM. Depending on the size of your installation, this can range from a single person to a complete team. The operations team is responsible to monitor, correct or work together with technical support to isolate any non-specific error condition. Basically, you have the following roles among your current or future ADSM operations team:

- **Operator:** Handles ADSM daily tasks and special server procedures.
- **Technical Support:** Installs new client code, upgrades server code, activates trace functions, deals with hardware configurations, and handles error conditions.
- **Storage Administrator:** Performs sanity checks in the environment, evaluates capacity planning issues, and manages storage pools.
- **System Administrator:** Manages ADSM policies, defines new nodes and management classes, and reports error conditions to technical support.

ADSM can provide a centralized point for backup, restore, archive, and retrieve operations. All of these functions require monitoring and special operational procedures. As a rule of thumb, consider having at least two people responsible for ADSM, so that each one can handle the workload in case of a problem (this way you can also avoid concentrating ADSM information in only one person). All the personnel directly involved in ADSM administration—the ADSM operations team—need to divide the ADSM workload into some basic activities:

- **Server procedures:** Start and stop the ADSM server.
- **Daily operations:** Check database free space, check activity log, check volumes in and out, send tapes to offsite, request to bring tapes onsite, check ADSM devices, back up ADSM database (or monitor its execution), and label new tapes.
- **Error or critical conditions:** Bring an offsite volume for restore, create another database or recovery log, deal with storage pool shortages, mirror write failures, communication errors, device problems, media errors, offsite tape not found, and any other non-specific problems.

12.1 Server Procedures

In this section we show how to start and stop an ADSM server. To a great extent, ADSM is a software program that runs as does any other application in your system. This means that you can start the server and let it run continuously and uninterrupted for days, weeks, or even longer if desired.

12.1.1 Starting the ADSM Server

The ADSM server can run as either a foreground or a background process. The foreground method allows you to have an active open console, with all system messages displayed. From this server console, you are able to issue all administrative commands. The background method allows you to close the active

window without stopping the server process. This is the recommended way if you are planning to have the ADSM server automatically started or stopped.

This section explains how to start an ADSM server (we assume that it is not running yet). Otherwise, you must stop it first and issue the commands we explain. You can have the ADSM server automatically started, in which case you need to configure it as a background process. If you wish to have it started manually, then the ADSM server needs to be configured as a foreground process. For further details on how to start the ADSM server, refer to the ADSM Quick Start manual which comes with the product software.

12.1.1.1 Active Console Mode

To start an ADSM server in foreground, you can start the ADSM server by simply executing the DSMSERV program in the ADSM system directory. If you run this way, you will not be able to close the window, since all output messages will be redirected to standard output.

The next screen shows you how to manually start the ADSM server on AIX in active console mode. This server console session is established using a predefined administrator ID SERVER_CONSOLE. All ADSM system messages are displayed on the console. From this server console, you are able to issue all administrative commands. The ADSM server is ready to receive and send data to the client machines.

```

kindu:[/usr/lpp/adsmerv/bin]$ dsmserv

ANR7800I DSMSERV generated at 14:49:49 on Feb  4 1999.

ADSTAR Distributed Storage Manager for AIX-RS/6000
Version 3, Release 1, Level 2.15
Licensed Materials - Property of IBM

5765-C43 (C) Copyright IBM Corporation 1990, 1998. All rights reserved.
U.S. Government Users Restricted Rights - Use, duplication or disclosure
restricted by GSA ADP Schedule Contract with IBM Corporation.

ANR0900I Processing options file dsmserv.opt.
ANR0990I ADSM server restart-recovery in progress.
ANR0200I Recovery log assigned capacity is 8 megabytes.
ANR0201I Database assigned capacity is 12 megabytes.
ANR0306I Recovery log volume mount in progress.
ANR0353I Recovery log analysis pass in progress.
ANR0354I Recovery log redo pass in progress.
ANR0355I Recovery log undo pass in progress.
ANR0352I Transaction recovery complete.
ANR2100I Activity log process has started.
ANR1305I Disk volume /adsm/stgpool/dirs01 varied online.
ANR1305I Disk volume /adsm/stgpool/file01 varied online.
ANR0984I Process 1 for EXPIRATION started in the BACKGROUND at 15:23:38.
ANR2803I License manager started.
ANR0811I Inventory client file expiration started as process 1.
ANR8190I HTTP driver ready for connection with clients on port 1580.
ANR8200I TCP/IP driver ready for connection with clients on port 1500.
ANR2560I Schedule manager started.
ANR0812I Inventory file expiration process 1 completed: examined 0 objects,
deleting 0 backup objects, 0 archive objects, 0 DB backup volumes, and 0
recovery plan files. 0 errors were encountered.
ANR0985I Process 1 for EXPIRATION running in the BACKGROUND completed with
completion state SUCCESS at 15:23:38.
ANR0993I ADSM server initialization complete.
ANR2835I Server is licensed for 67 clients.
ANR2861I Server is licensed to support NETWORK connections.
ANR2856I Server is licensed to support Open Systems Environment clients.
ANR2855I Server is licensed to support space-managed clients.
ANR2860I Server is licensed to support disaster recovery manager.
ANR2863I Server is licensed to support Enterprise Administration functions.
ANR2864I Server is licensed to support Server-to-Server Virtual Volumes.
ANR2869I Server is licensed for Advanced Device Support.
adsm>
ANR8192I HTTPS driver ready for connection with clients on port 1543.

```

On Windows NT, you also can start the server in active console mode from the ADSM server utilities program. In this method, you will not be able to close your desktop, since ADSM is running in foreground mode. Figure 34 on page 198 shows the main ADSM server utilities screen and the Launch option to start ADSM in console mode.

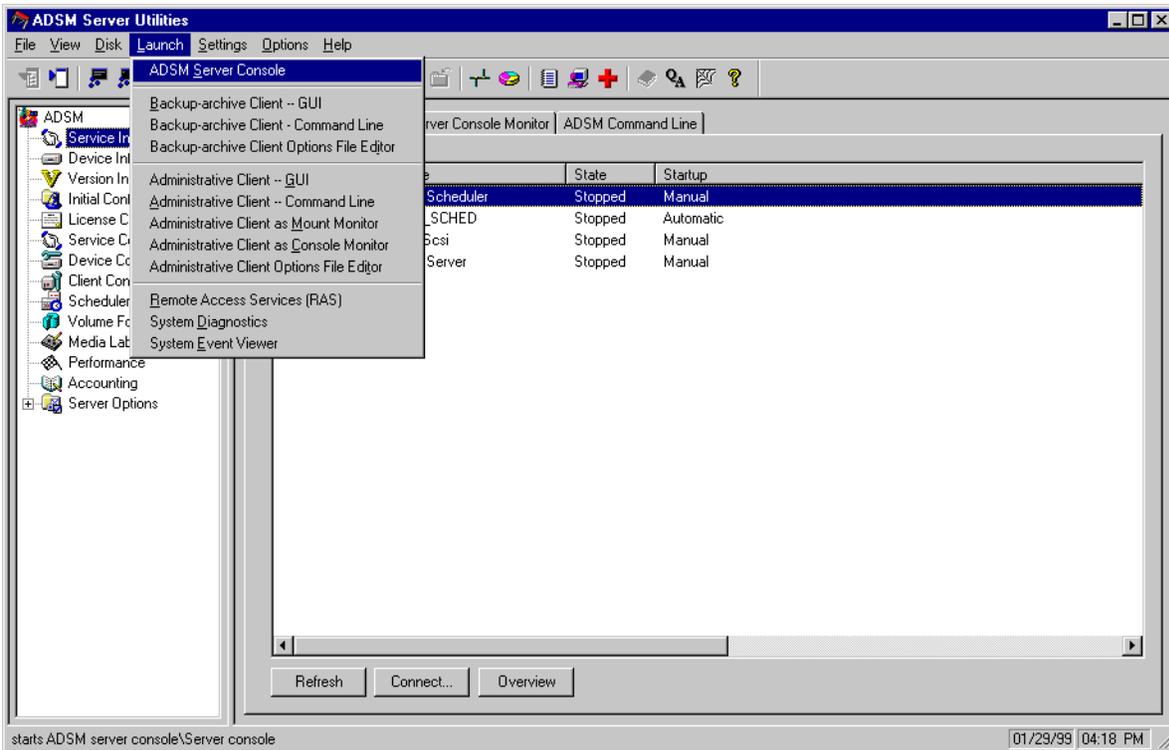


Figure 34. Launch Server Console in Foreground

By using the launch option, the ADSM console stays open as long as your desktop (login session) stays active. This is the recommended way if you plan to have an ADSM operation staff near the machine, or if you want to have all ADSM system activity displayed in the server, so that you can quickly check what is running.

12.1.1.2 Background Mode

You can start the ADSM server in background mode either manually or with each system start.

Manual Start

On AIX, to manually start the ADSM server in background mode, you have to run the script `rc.admserv`, which is a part of the ADSM server software.

```
kindu:[/usr/lpp/admserv/bin]$> nohup rc.admserv >/dev/console 2>&1 &
```

On Windows NT, the ADSM server is installed as a manual service. To start it, you have to go to Settings -> Control Panel -> Services, select the ADSM server service, and start it.

Automatic Start

On AIX, after the ADSM server installation, the system is configured to start the ADSM server automatically during the system startup. The system initialization file /etc/inittab is edited. The following entry is added to the file:

```
autosrvr:2:once:/usr/lpp/adsmerv/bin/rc.adsmerv >/dev/console 2>&1 #Start  
the ADSM server
```

On Windows NT, the ADSM server is installed as a manual service. To change its status to automatic, you have to go to Settings -> Control Panel -> Services, select the ADSM server service and change its startup setting to automatic.

The ADSM server will be started automatically the next time the system is rebooted. You can then check that the server is running by trying to connect to the server from an administrative client or a backup-archive client, or by checking whether the ADSM server process (AIX) or service (Windows NT) is running.

Administrative Console Session

If you start the ADSM server in the background, we highly recommend running an additional administrative client session in console mode on the administrative workstation. This session displays all ADSM messages, however, it is not an active session, that is, you cannot issue administrative commands here. This way you can easily monitor your ADSM system. The following command shows how to start an administrative session in console mode.

```
kindu:[/usr/lpp/adsm/bin]$> dsmdmc -id=sysadmin -passwd=philip -console  
ADSTAR Distributed Storage Manager  
Command Line Administrative Interface - Version 3, Release 1, Level 0.6  
(C) Copyright IBM Corporation, 1990, 1997, All Rights Reserved.  
  
Session established with server AIX1: AIX-RS/6000  
Server Version 3, Release 1, Level 2.15  
Server date/time: 03/17/1999 23:44:24 Last access: 03/17/1999 23:43:08
```

12.1.2 Stopping the ADSM Server

There are two ways of stopping an ADSM server:

- Immediate stop: This brings the server down immediately. All system processes and client activity are interrupted.
- Drain system activity and stop: This is the recommended way, because you can be confident about the processes that are running and what you may need to restart later on. Performing basic pre-checks and then stopping gives you (and the operations team) an idea of the workload that may need to be restarted.

12.1.2.1 Stopping the ADSM Server Immediately

You can halt the server without warning if an unplanned operating system problem requires the server to be stopped.

On Windows NT, you can terminate a server session by stopping the service. To change the status of the server service, you have to go to Settings -> Control Panel -> Services, select the ADSM server service and stop it.

When you halt the server, all processes are abruptly stopped and client sessions are canceled, even if they are not completed. Any in-progress transactions are

rolled back when the server is restarted. When the server is halted, neither administrator activity nor client operations are possible. We recommend that you halt the server only after current administrative and client node sessions have completed or canceled.

You must use the `halt` command to shutdown the ADSM server:

```
adsm>halt
ANR2017I Administrator SERVER_CONSOLE issued command: HALT
ANR7835I ADSM thread 1 terminated in response to server shutdown.
ANR7835I ADSM thread 10 terminated in response to server shutdown.
:
:
ANR7835I ADSM thread 42 terminated in response to server shutdown.
ANR7835I ADSM thread 43 terminated in response to server shutdown.
ANR0991I ADSM server shutdown complete.
```

Note:

The QUIESCE option with the `halt` command is recommended only if you plan to do a database dump by using the `DSMSERV DUMPDB` command immediately after halting. Because ADSM supports online database backup (`backup db` command), the `DSMSERV DUMPDB` command should rarely, if ever, be needed.

12.1.2.2 Draining System Activity before Stopping the ADSM Server

To shut down the server without severely impacting administrative and client node activity with the server, we recommend that you take some steps before actually issuing the `halt` command to the server:

1. Disable the server: This prevents new client node sessions from starting. Disabling new client sessions prevents users from establishing client node sessions with the server. This command does not affect current client sessions in progress or system processes like migration and reclamation. To disable client node access to the server, use the `disable sessions` command:

```
adsm> disable sessions
ANR2097I Server now disabled for backup/archive client access.
```

When you disable client sessions from 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.

You can issue the `query status` command to determine if the server is enabled or disabled.

2. Query for session information: This is to identify any existing administrative and client node sessions.

When administrators or clients access ADSM, either 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 the `query session` command:

```
adsm> query session
```

Sess Number	Comm. Method	Sess State	Wait Time	Bytes Sent	Bytes Recvd	Sess Type	Platform	Client Name
1	Tcp/Ip	Run	0 S	800	170	Admin	WinNT	ADMIN
2	Tcp/Ip	IdleW	3.4 M	715	194	Node	WinNT	WKS.23FFHBV
3	Tcp/Ip	IdleW	7 S	705	186	Node	AIX	AIXCLIENT1
22	Tcp/Ip	IdleW	3 S	1.1 K	211	Node	AIX	DEVNODE1

3. Notify users: This is to 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: To cancel a session, you must identify it by its session number. You can identify the session number by issuing the `query session` command. For example, in the screen above, one session number is 2 (client WKS.23FFHBV). You can cancel that session by entering the `cancel session` command:

```
adsm> cancel session 2
```

```
ANR0490I Canceling session 2 for node WKS.23FFHBV (WinNT) .
```

If you want to cancel all backup-archive client sessions, enter the following command:

```
adsm> cancel session all
```

```
ANR0490I Canceling session 3 for node AIXCLIENT1 (AIX) .  
ANR0490I Canceling session 22 for node DEVNODE1 (AIX) .
```

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 that are not yet committed to the database are undone. If necessary, the cancellation process may be delayed. For example, if a client restore session is in a restartable state, the file space is locked and no files can be moved from sequential volumes. This prevents the data from being migrated, moved, reclaimed, or backed up by another operation. These sessions will automatically expire when the specified restore interval has passed. For further details, see the `query restore` and `cancel restore` commands.

5. Find out if any other processes are running, such as server migration or inventory expiration, by using the `query process` command. If a database backup process is running, allow it to complete before halting the server. If other types of processes are running, cancel them by using the `cancel process` command.
6. Halt the server and all server operations by using the `halt` command as explained in Chapter 12.1.2.1, “Stopping the ADSM Server Immediately” on page 199.

12.2 Event Monitoring

One of the features that are integrated to ADSM is a complete set of messages that can be enabled, disabled, stored, or even forwarded to other ADSM servers, user programs, enterprise consoles like T/EC, or using SNMP. These components are also known as receivers. Since event logging has a full set of functions and features that are beyond of the scope of this book, we only cover some basic settings.

ADSM divides the range and the source of the messages by prefixes. For example, prefix ANR is for server messages, and most of prefix ANS messages are from backup-archive client sessions. Details of all these messages are explained in *ADSM V3R1 Messages*, (GC35-0271).

You can enable or disable messages for any particular receiver. The only restriction is with respect to severe or internal system messages that will not be disabled, even if you define it.

Logging of events to both the ADSM server console and the ADSM activity log begins automatically at server startup. Messages can appear on the server console, the administrative client, an operator terminal, the administrative graphical user interface, the backup-archive client, or the space-management client.

ADSM provides an activity log to help the administrator track server activity and monitor the system. The activity log contains messages generated by the server and is stored in the database. Any messages sent to the server console are stored in the activity log. Examples of the types of messages stored in the activity log include:

- When client sessions start or end
- When migration starts or ends
- When backed up files are expired from data storage
- Any output generated from background processes

You can check if event logging is enabled for a receiver by using the `query status` command and checking the Active Receivers field:

```

adsm> query status
ADSM Server for AIX-RS/6000 - Version 3, Release 1, Level 2.1

Server Name: AIX1
Server host name or IP address:
Server TCP/IP port number:
Server URL:
Crossdefine: Off
Server Password Set: No
Server Installation Date/Time: 02/13/1998 10:59:41
Server Restart Date/Time: 02/02/1999 11:14:31
Authentication: On
Password Expiration Period: 120 Day(s)
Invalid Sign-on Attempt Limit: 3
Minimum Password Length: 0
WEB Admin Authentication Time-out (minutes): 0
Registration: Open
Availability: Enabled
Accounting: On
Activity Log Retention Period: 1 Day(s)
License Audit Period: 30 Day(s)
Last License Audit: 01/23/1999 01:58:01
Server License Compliance: Valid
Central Scheduler: Active
Maximum Sessions: 25
Maximum Scheduled Sessions: 12
Event Record Retention Period: 10 Day(s)
Client Action Duration: 5 Day(s)
Schedule Randomization Percentage: 25
Query Schedule Period: Client's Choice
Maximum Command Retries: Client's Choice
Retry Period: Client's Choice
Scheduling Modes: Any
Log Mode: Normal
Database Backup Trigger: Disabled
Active Receivers: CONSOLE ACTLOG
Configuration manager?: Off
Refresh interval: 60
Last refresh date/time:

```

We recommend that you run the `begin eventlogging` command to make sure that all defined receivers are active. Use this command to begin logging events to one or more receivers. Event logging automatically begins when the server is started for the console and activity log, and for any receivers that are started automatically based on entries in the server options file. A receiver for which event logging has begun is an active receiver. The next screen shows you how to activate all defined receivers:

```

adsm> begin eventlogging

ANR1825I Event logging active for the CONSOLE receiver.
ANR1825I Event logging active for the ACTLOG receiver.

```

We recommend that you enable nodename logging, so that you can also monitor client activity centrally in the server. Note that enabling client events to the activity log will increase the ADSM database utilization. You can set a retention period for the log records by using the `set actlogretention` command. To enable a specific receiver (CONSOLE, ACTLOG, EVENTSERVER, FILE, SNMP, TIVOLI, USEREXIT) enter the `enable events` command:

```

adsm> enable events actlog all nodename=*

ANR1844I ENABLE EVENTS command processed.

```

Windows NT Users:

On Windows NT, you also have a NTEVENTLOG receiver. This enables you to send events to the application eventlog, so that you can use the Windows Event Viewer to monitor an ADSM system. In this case, we recommend that you issue the `enable events` command for the events category SEVERE only.

12.3 Daily Sanity Checks

This section shows you how to perform some basic, yet important, system validations on your ADSM environment. These checks assume that you do not have any kind of automation in place, so that you (or the ADSM operations team) may have to control some of the ADSM aspects to have it running properly.

All commands are based in the administrative command line interface. This makes the monitoring process easier and can be further enhanced to accommodate any automation facility you may desire. We recommend that you start using all of the commands on a daily basis, then once you feel confident of your environment settings, you can use any or all of them on a weekly basis. You should balance the level of information that you believe is necessary as you start running the commands.

You should run monitoring commands as part of the daily operations check, as a sync-point for all ADSM personnel, especially if your company has different working hours and different working teams.

ADSM has a rich set of commands that allow you to monitor and create your own reports. The `query` command set is the base for most information from the server. Typing `help query` or `help query <command>` provides you with further details on all possible combinations.

12.3.1 Database and Recovery Log

The following is a list of important ADSM queries for the database and recovery log in alphabetical order.

12.3.1.1 Display Database Information

To display information about your current database occupancy and how much space is still available for extension, use the `query db` command:

```

adsm> query db

Available Space (MB)  Assigned Capacity (MB)  Maximum Extension (MB)  Maximum Reduction (MB)  Page Size (bytes)  Total Usable Pages  Used Pages  Pct Util  Max. Pct Util
-----
                20                20                0                12            4,096            5,120            772      15.1      15.1

adsm> query db format-detail

Available Space (MB): 20
Assigned Capacity (MB): 20
Maximum Extension (MB): 0
Maximum Reduction (MB): 12
Page Size (bytes): 4,096
Total Usable Pages: 5,120
Used Pages: 772
Pct Util: 15.1
Max. Pct Util: 15.1
Physical Volumes: 6
Buffer Pool Pages: 512
Total Buffer Requests: 18,350
Cache Hit Pct.: 98.70
Cache Wait Pct.: 0.00
Backup in Progress?: No
Type of Backup In Progress:
Incrementals Since Last Full: 0
Changed Since Last Backup (MB): 3.02
Percentage Changed: 100.00
Last Complete Backup Date/Time: 01/27/1999 15:58:05

```

12.3.1.2 Display Database Volumes Information

Use the `query dbvolume` command to display information on one or more database volumes, including available, allocated, and free space on the volume. This command displays information about the specified database volume and any database volume copies.

```

adsm> query dbvolume

Volume Name      Copy      Volume Name      Copy      Volume Name      Copy
(Copy 1)         Status    (Copy 2)         Status    (Copy 3)         Status
-----
/adsm/database/- Sync'd  /adsm/database/- Sync'd
primary/file01   copy/file01
/usr/lpp/adsmse-  Sync'd
rv/bin/db.dsm    Undef-
                  ined
                  Undef-
                  ined

```

12.3.1.3 Display Information on the Recovery Log

Use the `query log` command to display allocation information about the recovery log, including utilization, expansion, and reduction abilities.

```

adsm> query log

Available Assigned Maximum Maximum Page Total Used Pct Max.
Space Capacity Extension Reduction Size Usable Pct Util Pct
(MB) (MB) (MB) (MB) (bytes) Pages Pages Util Pct
-----
      8      8      0      4  4,096  1,536   268  17.4  30.5

adsm> query log format=detail

Available Space (MB): 8
Assigned Capacity (MB): 8
Maximum Extension (MB): 0
Maximum Reduction (MB): 4
Page Size (bytes): 4,096
Total Usable Pages: 1,536
Used Pages: 268
Pct Util: 17.4
Max. Pct Util: 30.5
Physical Volumes: 1
Log Pool Pages: 128
Log Pool Pct. Util: 0.19
Log Pool Pct. Wait: 0.00
Cumulative Consumption (MB): 15.92
Consumption Reset Date/Time: 10/29/1998 12:03:23

```

12.3.1.4 Information on Log Volumes

Use the `query logvolume` command to display information on one or more recovery log volumes, including available, allocated, and free space on the volume. This command displays information about the specified recovery log volume and any recovery log volume copies. The next screen shows you the output for a UNIX server.

```

adsm> query logvolume

Volume Name      Copy   Volume Name      Copy   Volume Name      Copy
(Copy 1)         Status (Copy 2)         Status (Copy 3)         Status
-----
/adsm/log/prim-  Sync'd /adsm/log/copy/-  Sync'd              Undef-
ary/file01      file01
/usr/lpp/adsmse-  Sync'd              Undef-              Undef-
rv/bin/log.dsm  ined                ined                 ined

```

12.3.1.5 Display Sequential Volume History Information

Use the `query volhistory` command to display sequential volume history information that has been collected by the server.

```
adsm> query volhist begindate=02/04/1999
```

```
    Date/Time: 02/04/1999 15:13:11
    Volume Type: STGDELETE
    Backup Series: 0
    Backup Operation: 0
    Volume Seq: 0
    Device Class: CDLT
    Volume Name: VOL004
    Volume Location:
```

```
    Date/Time: 02/05/1999 16:55:44
    Volume Type: STGNEW
    Backup Series: 0
    Backup Operation: 0
    Volume Seq: 0
    Device Class: CDLT
    Volume Name: V00001
    Volume Location:
```

12.3.2 Data Storage

The following is a list of important ADSM queries for the data storage in alphabetical order.

12.3.2.1 Client Node Storage Utilization

Use the `query auditoccupancy` command to display information about the client node server storage utilization. The displayed information is current as of the last license audit processed by the ADSM server. You can use this information to determine if and where client node storage utilization needs to be balanced. This information can also assist with billing clients for storage usage.

```
adsm> query auditoccupancy
```

```
License information as of last audit on 02/05/99 at 17:20:43.
```

Node Name	Backup Storage Used (MB)	Archive Storage Used (MB)	Space-Managed Storage Used (MB)	Total Storage Used (MB)
NOGALES	0	0	0	0
SQLCLIENT	0	0	0	0
BRUNO	0	0	0	0
SCANDIUM_EXCH	0	0	0	0
NOGALES_EXCH	0	0	0	0
SCANDIUM	0	619	0	619
PAGOPAGO	0	0	0	0
THORIUM	0	0	0	0
PAGOPAGO_EXCH	44	0	0	44
LADOGA	0	0	0	0
SNOWHILL	0	0	0	0
WKS.23FFHEV	101	7	0	108
23FFHEV	0	0	0	0
SRV.ITSOALO	0	0	0	0
AIXCLIENT1	0	0	0	0
MAINPRD0	0	0	0	0
VANSTOR	0	0	0	0

12.3.2.2 Display Drive Information

Use the `query drive` command to display information about a drive located in a server-attached library. You may use this command to validate if all drives are online. The next screen shows you the output for a UNIX server.

```
adsm> query drive
```

Library Name	Drive Name	Device Type	Device	ON LINE
L8MM	DRV0	8MM	/dev/mt0	Yes
LDLT	DRV1	DLT	/dev/mt1	Yes

12.3.2.3 Query a Library Volume

Use the `query libvolume` command to display information about one or more volumes that have been previously checked into an automated library for use by the AD SM server.

```
adsm> query libvolume
```

Library Name	Volume Name	Status	Last Use	Home Element
LDLT	V00001	Scratch		8
LDLT	V00002	Scratch		9

12.3.2.4 Sequential Access Storage Pool Media

Use the `query media` command to display information about the sequential access primary and copy storage pool volumes moved by the `move media` command.

```
adsm>query media * stg=offdata
```

```
ANR2017I Administrator SERVER_CONSOLE issued command: QUERY MEDIA * stg=offdata
```

Volume Name	State	Location	Automated LibName
V00003	Mountable not in library	ITSO vault	

12.3.2.5 Display File Space Information by Storage Pool

Use the `query occupancy` command to display information on where a client's file spaces are stored and how much space they occupy. This command can be used to display information about file spaces that are stored in primary or copy storage pools.

```

adsm> query occupancy

```

Node Name	Type	Filespace Name	Storage Pool Name	Number of Files	Physical Space Occupied (MB)	Logical Space Occupied (MB)
AIXCLIENT1	Bkup	/	BACKUPPOOL	4	0.01	0.01
AIXCLIENT1	Bkup	/home	BACKUPPOOL	13	0.17	0.17
AIXCLIENT1	Arch	/home	BACKUPPOOL	13	0.17	0.17
PAGOPAGO_EXCH	Bkup	ACME MOTOR CORP.WES-T.PAGOPA-GO.IS	8MMPool	1	10.00	10.00
PAGOPAGO_EXCH	Bkup	ACME MOTOR CORP.WES-T.PAGOPA-GO.DIR	8MMPool	5	34.04	34.04
SCANDIUM	Arch	CDRIVE	BACKUPPOOL	251	619.05	619.05
WKS.23FFHBV	Bkup	\\23ffhbv-c\$	BACKUPPOOL	1,536	100.93	100.93
WKS.23FFHBV	Arch	\\23ffhbv-c\$	BACKUPPOOL	146	6.52	6.52

12.3.2.6 Query One or More Pending Mount Requests

Use the `query request` command to display information about one or more pending mount requests.

```

adsm>query request
ANR2017I Administrator SERVER_CONSOLE issued command: QUERY REQUEST
ANR8352I Requests outstanding:
ANR8306I 002: Insert DLT volume ML2611 R/W into the slot with element number 3
of library LDLT within 59 minutes; issue 'REPLY' along with the request ID when
ready.

```

12.3.2.7 Query One or More Storage Pools

Use the `query stgpool` command to display information about one or more storage pools.

```

adsm> query stgpool

```

Storage Pool Name	Device Class Name	Estimated Capacity (MB)	Pct Util	Pct Migr	High Mig Pct	Low Mig Pct	Next Storage Pool
8MMPool	8MMDEV	4,944.0	0.9	100.0	90	70	
ARCHIVEPOOL	DISK	0.0	0.0	0.0	90	70	BACKUPPOOL
BACKUPPOOL	DISK	736.0	75.5	75.5	80	70	
DISKDATA	DISK	50.0	0.0	0.0	70	30	TAPEDATA
DISKDIRS	DISK	50.0	0.0	0.0	100	70	
OFFDATA	COFFSITE	0.0	0.0				
OFFDIRS	COFFSITE	0.0	0.0				
SPACEMPool	DISK	0.0	0.0	0.0	90	70	
TAPEDATA	CDLT	204,800,00	0.0	0.0	100	70	

12.3.2.8 Query One or More Storage Pool Volumes

Use the `query volume` command to display information about one or more storage pool volumes. This command displays information about volumes defined to ADSM. The next screen shows you the output for a UNIX server.

```
adsm> query volume
```

Volume Name	Storage Pool Name	Device Class Name	Estimated Capacity (MB)	Pct Util	Volume Status
/dsm_stg1/data1.dsm	BACKUPPOOL	DISK	368.0	54.3	On-Line
/dsm_stg2/data2.dsm	BACKUPPOOL	DISK	368.0	96.8	On-Line
/dsm_stg3/dirsdata1.dsm	DISKDIRS	DISK	50.0	0.0	On-Line
/dsm_stg3/diskdata1.dsm	DISKDATA	DISK	50.0	0.0	On-Line
ADSM01	8MMPool	8MDEV	4,944.0	0.9	Filling
V00001	TAPEDATA	CDLT	20,480.0	0.8	Filling

12.3.3 Client-Server Activity

The following is list of important AD SM queries for client-server activities in alphabetical order.

12.3.3.1 Search Activity Log for Messages

One of the basic commands to access the activity log is the `query actlog` command. Use the `query actlog` command to search the server activity log for messages selected by a range of time or date, message number, or string expression. If you do not specify any parameters for this command, all messages that were generated in the last hour are displayed.

```
adsm> query actlog
```

```
02/02/1999 15:35:01 ANR2561I Schedule prompter contacting SCANDIUM (session
25) to start a scheduled operation.
02/02/1999 15:35:03 ANR0406I Session 26 started for node SCANDIUM (WinNT)
(Tcp/Ip 9.1.150.184(3244)).
02/02/1999 15:35:26 ANR0406I Session 27 started for node SCANDIUM (WinNT)
(Tcp/Ip 9.1.150.184(3247)).
...
02/02/1999 15:35:28 ANE4967I (Session: 27, Node: SCANDIUM) Aggregate data
transfer rate: 0.81 KB/sec
02/02/1999 15:35:28 ANE4968I (Session: 27, Node: SCANDIUM) Objects compressed
by: 0%
02/02/1999 15:35:28 ANE4964I (Session: 27, Node: SCANDIUM) Elapsed processing
time: 00:00:01
02/02/1999 15:35:28 ANR0403I Session 27 ended for node SCANDIUM (WinNT).
02/02/1999 15:35:28 ANR0406I Session 28 started for node SCANDIUM (WinNT)
(Tcp/Ip 9.1.150.184(3249)).
02/02/1999 15:35:33 ANR0522W Transaction failed for session 28 for node
SCANDIUM (WinNT) - no space available in storage pool
BACKUPPOOL and all successor pools.
02/02/1999 15:35:41 ANR0403I Session 28 ended for node SCANDIUM (WinNT).
02/02/1999 15:36:21 ANR0403I Session 26 ended for node SCANDIUM (WinNT).
02/02/1999 15:36:21 ANR0406I Session 29 started for node SCANDIUM (WinNT)
(Tcp/Ip 9.1.150.184(3274)).
02/02/1999 15:36:21 ANR0403I Session 29 ended for node SCANDIUM (WinNT).
02/02/1999 15:50:09 ANR0984I Process 2 for MIGRATION started in the BACKGROUND
at 15:50:09.
02/02/1999 15:50:09 ANR1000I Migration process 2 started for storage pool
BACKUPPOOL.
02/02/1999 15:50:09 ANR1025W Migration process 2 terminated for storage pool
BACKUPPOOL - insufficient space in subordinate storage
pool.
02/02/1999 15:50:09 ANR0985I Process 2 for MIGRATION running in the BACKGROUND
completed with completion state FAILURE at 15:50:09.
02/02/1999 15:50:09 ANR1002I Migration for storage pool BACKUPPOOL will be
retried in 60 seconds.
```

12.3.3.2 Query Scheduled and Completed Events

Use the `query event` command to display scheduled and completed events. This command takes two forms, depending on whether the query applies to scheduled client operations or scheduled administrative commands.

Each scheduled client operation and administrative command is called an event. The server tracks each scheduled event and records the results of each function in the database. An event record is created whenever processing of a scheduled command is started or missed.

The `query event` command may display information about an event for which there is no event record. For example, if you request information about a scheduled event in the future, the `query event` command displays the status of Future for the event even though no event record has been created for this event. In addition, the `query event` command displays the status for past events whose event records have already been removed from the database. In this case, the server returns the status of Uncertain because the actual event status cannot be determined without the event records. The following shows the command and its output to query scheduled client events:

```
adsm> query event * *  
  
Scheduled Start      Actual Start          Schedule Name Node Name      Status  
-----  
02/10/1999 00:00:00 02/10/1999 00:00:04 EXCH_DAILY      SCANDIUM_EXCH Completed  
02/10/1999 00:05:00                MYSTUFF         NOGALES        Missed  
02/10/1999 00:35:00 02/10/1999 00:35:03 HOURLY_OFFLI-  SCANDIUM        Completed  
                NE_ARCH  
02/10/1999 01:05:00                MYSTUFF         NOGALES        Missed  
02/10/1999 01:35:00 02/10/1999 01:35:02 HOURLY_OFFLI-  SCANDIUM        Completed  
                NE_ARCH  
02/10/1999 02:05:00                MYSTUFF         NOGALES        Missed  
02/10/1999 02:35:00 02/10/1999 02:35:03 HOURLY_OFFLI-  SCANDIUM        Completed  
                NE_ARCH  
02/10/1999 03:05:00                MYSTUFF         NOGALES        Missed  
02/10/1999 03:35:00 02/10/1999 03:35:03 HOURLY_OFFLI-  SCANDIUM        Completed  
                NE_ARCH  
02/10/1999 04:05:00                MYSTUFF         NOGALES        Missed  
02/10/1999 04:35:00 02/10/1999 04:35:03 HOURLY_OFFLI-  SCANDIUM        Completed  
                NE_ARCH  
02/10/1999 05:05:00                MYSTUFF         NOGALES        Missed  
02/10/1999 05:35:00 02/10/1999 05:35:03 HOURLY_OFFLI-  SCANDIUM        Completed  
                NE_ARCH  
02/10/1999 06:05:00                MYSTUFF         NOGALES        Missed
```

The following shows the command and its output to query scheduled administrative events:

```
adsm> query event * type=administrative
```

Scheduled Start	Actual Start	Schedule Name	Status
03/16/1999 00:00:00	03/16/1999 00:00:29	AUDIT_LICENSE	Completed
03/16/1999 04:00:00	03/16/1999 04:00:29	REDBOOK_OFFSITE	Completed
03/16/1999 07:00:00	03/16/1999 07:00:00	BACKUP_DATABASE	Failed
03/16/1999 07:00:00	03/16/1999 07:00:00	DELETE_VOLHIST	Completed
03/16/1999 07:00:00	03/16/1999 07:00:00	MIGRATION_START	Completed
03/16/1999 10:00:00	03/16/1999 10:00:00	MIGRATION_STOP	Completed
03/16/1999 10:00:00	03/16/1999 10:00:00	RECLAIM_OFFDERS_START	Completed
03/16/1999 11:00:00	03/16/1999 11:00:00	RECLAIM_OFFDATA_START	Completed
03/16/1999 11:00:00	03/16/1999 11:00:00	RECLAIM_OFFDERS_STOP	Completed
03/16/1999 14:00:00	03/16/1999 14:00:01	RECLAIM_OFFDATA_STOP	Completed

12.3.3.3 Query One or More Server Processes

Use the query process `query process` command to display information about one or more active background processes.

```
adsm> query process
```

Process Number	Process Description	Status
24	Move Data	Volume /dsm_stg1/data1.dsm (storage pool BACKUPPOOL), Target Pool TAPEDATA, Moved Files: 0, Moved Bytes: 0, Unreadable Files: 0, Unreadable Bytes: 0. Current Physical File (bytes): 4,096 Waiting for mount of scratch volume (72 seconds).
25	Database Backup	Full backup: 0 pages of 1010 backed up.

12.3.3.4 Query Restartable Restore Sessions

Use the `query restore` command to display information about the restartable restore sessions recorded in the server database. Certain restore operations invoke a special protocol with ADSM. These special restores are called restartable restore sessions. It is desirable to know which restores are using this protocol for two reasons:

1. Restartable restore sessions that fail for some reasons (network outage, client failure, or server outage) can be restarted from where the restore operation left off. This can save valuable time because these restore operations do not have to be started over again. The `query restore` command displays the restartable restore sessions.
2. Restartable restore sessions lock the filesystem, and do not allow files to be moved off of sequential volumes. The `query restore` command displays the restartable restore sessions and their associated filesystem. The `cancel restore` command can be used to cancel a restartable restore session.

```

adsm> query restore

  Sess   Restore   Elapsed   Node Name   Filespace
Number  State      Minutes
-----
  -1     Restartable  2         WKS.23FFHEV  \\michigan-
                                         \adsmteam

```

12.3.3.5 Query One or More Client Sessions

Use the `query session` command to display information on one or more administrative and client node sessions.

```

adsm> query session

  Sess  Comm.  Sess   Wait   Bytes   Bytes  Sess  Platform  Client Name
Number Method State   Time   Sent   Recvd  Type
-----
  258   Tcp/Ip  Run    0 S    80.4 K  1.8 K  Admin  WinNT     ARMANDO
  265   Tcp/Ip  IdleW  18 S    826     186   Node   AIX       AIXCLIENT1
  266   Tcp/Ip  IdleW  7 S    828     194   Node   WinNT     WKS.23FFHEV

```

12.4 Tape Management

This section shows you how to use ADSTM tapes, how to label them, and how to manage the many possible stages a tape can have.

Tape is a vital component of the storage hierarchy. When you start using ADSTM, you have to create storage pools to hold all your data. After a storage pool is defined, you also need to define volumes to the storage pool, so that you have space to store the data. ADSTM allows you to use and reuse removable media to store data. One of the key concepts for ADSTM is the distinction between onsite data and offsite copy. For our discussion, we assume that all data that is available in the company is onsite data. All data that is remotely stored in another location is offsite data. Figure 35 on page 213 shows you onsite and offsite tapes.

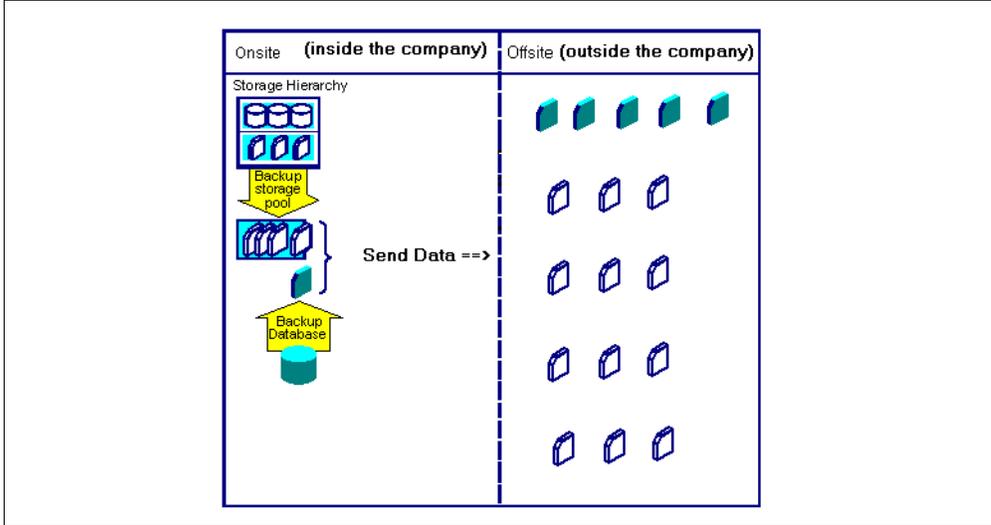


Figure 35. Onsite and Offsite Distinction

All tapes used for onsite data are copied to different tapes to be sent to an offsite location. ADSM keeps track of all volumes (onsite and offsite) so that you do not have to care about which volumes are in the company or in the offsite location. ADSM can handle both onsite and offsite tapes. Figure 36 on page 214 shows a generic tape processing from the moment the data is backed up until it is sent to offsite.

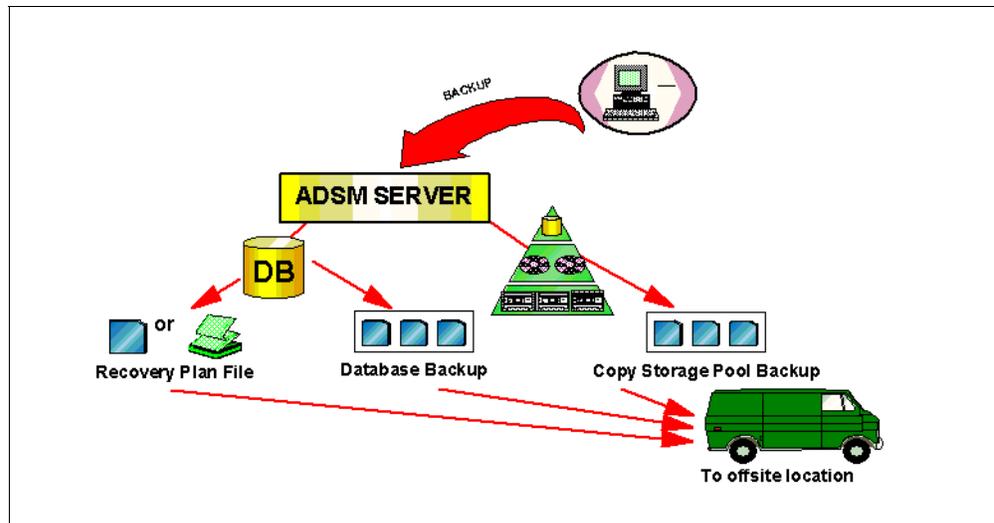


Figure 36. ADSM Tape Processing

Under normal conditions, all offsite tapes contain exactly the same valid data that you have onsite. This means that onsite and offsite volumes mirror the same data. This does not mean that you have the same amount of tapes, but simply shows that what you have “inside” is the same that you have “outside” (regardless of the number of tapes).

Depending on the environment you are working on (manual devices or automated libraries), you may have different operational procedures. For example, if you have a 3494 tape library, the CHECKIN and CHECKOUT operations are part of the tape movement, with no manual intervention required. Otherwise, if you have a single I/O slot in a library (or even a manual device), you need to REPLY to the mount/unmount requests before executing the next command.

As you can see in Figure 37 on page 215, the steps that you must accomplish to store data on tape are:

1. Label tape (create volume)
2. Add tape volume to ADSM inventory
3. Use the tapes (ADSM will do this)
4. Create Offsite copies of your data
5. Send Offsite volumes to safe location
6. Manage expired data (reclamation)
7. Receive free Offsite volumes no longer needed (reclaimed)

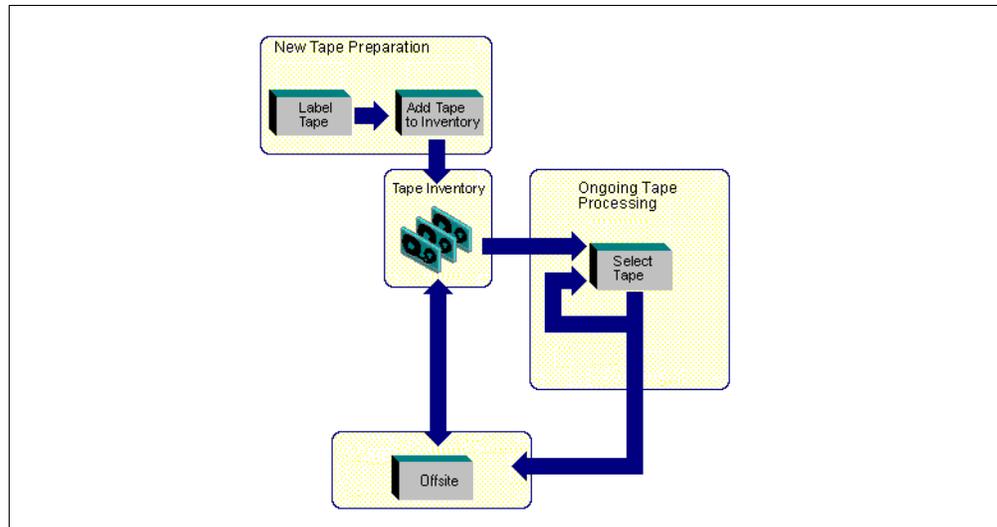


Figure 37. How ADSM Tapes Are Processed

Although it is not in the scope of this book to show you how to use DRM, it is a key component to successfully handle many tape management procedures. We recommend that you further evaluate DRM as a better approach to handle offsite tapes, database backups, and disaster plans. As an example, Figure 38 on page 215 shows a typical DRM tape management with enhanced control of the actual state of the tapes:

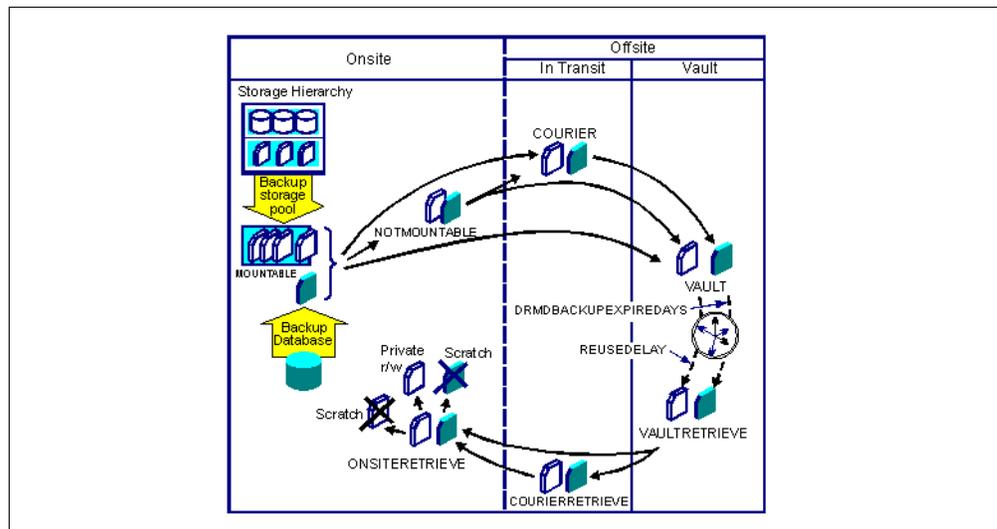


Figure 38. ADSM DRM Tape Lifecycle

12.4.1 Labelling Tapes

To start using new tapes, you need to first prepare removable media by using the `label libvolume` command. If you have an automated library (like the example we show), you can load the library with the new volumes and label them all with one single command. Depending on the library that you have (3494, SCSI, or

MANUAL) the command syntax may be slight different. For further details, type `help label libvolume`.

To label new tapes (scratch tapes) in a SCSI library which has no barcode reader, use the following command:

```
adsm>label libvolume ldlt search=yes labelsorce=prompt checkin=scratch

ANR2017I Administrator SERVER_CONSOLE issued command: LABEL libv ldlt
search=yes labelsorce=prompt checkin=scratch
ANR0984I Process 1 for LABEL LIBVOLUME started in the BACKGROUND at 11:01:42.
ANR8799I LABEL LIBVOLUME: Operation for library LDLT started as process 1.
ANR8809I 001: Please provide the label name for the volume in slot element 5 of
library LDLT by issuing REPLY n LABEL=xxx within 60 minutes, where n is the
request ID and xxx is the desired label name.

adsm>reply 1 label=CD2201

ANR2017I Administrator SERVER_CONSOLE issued command: REPLY 1 label=CD2201
ANR8499I Command accepted.
ANR8810I Volume CD2201 has been labeled in library LDLT.
ANR8427I CHECKIN LIBVOLUME for volume CD2201 in library LDLT completed
successfully.
ANR8809I 002: Please provide the label name for the volume in slot element 9 of
library LDLT by issuing REPLY n LABEL=xxx within 60 minutes, where n is the
request ID and xxx is the desired label name.

adsm>reply 2 label=vol002

ANR2017I Administrator SERVER_CONSOLE issued command: REPLY 2 label=vol002
ANR8499I Command accepted.
ANR8810I Volume VOL002 has been labeled in library LDLT.
ANR8427I CHECKIN LIBVOLUME for volume VOL002 in library LDLT completed
successfully.
ANR8801I LABEL LIBVOLUME process 1 for library LDLT completed; 2 volume(s)
labelled, 2 volume(s) checked-in.
ANR0985I Process 1 for LABEL LIBVOLUME running in the BACKGROUND completed with
completion state SUCCESS at 11:06:44.
```

If you have a library of type MANUAL (for example, a library named `l8mm`) and you want to label a volume as `VOL001`, use the following command:

```
adsm> label libvolume l8mm vol001 overwrite=yes
ANS8003I Process number 23 started.

adsm> query process

  Process Process Description  Status
  Number
-----
      23 LABEL LIBVOLUME      ANR8804I Labelling volume VOL001 in library
                                L8MM.

adsm> query actlog search=vol001

02/05/1999 15:54:30 ANR2017I Administrator ADMIN issued command: LABEL libv
                                l8mm vol001 overwrite=yes
02/05/1999 15:54:30 ANR8326I 017: Mount 8MM volume VOL001 R/W in drive DRVO
                                (/dev/mt0) of library L8MM within 60 minutes.
02/05/1999 15:56:48 ANR8372I 017: Remove 8MM volume VOL001 from drive DRVO
                                (/dev/mt0) of library L8MM.
02/05/1999 15:56:48 ANR8800I LABEL LIBVOLUME for volume VOL001 in library L8MM
                                completed successfully.
```

12.4.2 Onsite Tape Management to Offsite

If you think of onsite tape management as the source of all data, then all commands that you run from this point on are to send data to an offsite location, which can be thought as a destination. The commands explained in this section shows how to move data from onsite to offsite.

For manually mounted devices, ADSM sends messages that request that volumes be mounted when they are needed. When ADSM needs to write to a tape, it first tries to write to a previously used tape with enough room for new data. If that is not possible, ADSM asks to mount a SCRATCH (new) tape for use. In manual libraries, you do not need to incorporate them into a library since there is no inventory.

For devices in automated libraries (such as a tape autochanger), ADSM interacts with the library to mount volumes, but sends messages when the library needs attention from an operator. ADSM also tracks the inventory of media in each automated library. You need to manage new tapes that you have labelled by the `label libvolume` command or they are reclaimed volumes from offsite.

The steps that you need to perform to send data to offsite are the following:

1. Backup your vital data to the copy storage pools
2. List all volumes that must be sent to offsite
3. Check out volumes from the library
4. Update the volume location to offsite

12.4.2.1 Backup Your Vital Data

Use the `backup stgpool` command to create copies of files that reside in a primary storage pool, and store them in a copy storage pool. If a file is already duplicated in the specified copy storage pool, a new copy of the file is not made in that copy pool. However, if a copy of a physical file already exists in the copy storage pool, but the copy is marked damaged, a new copy will be created, provided that the primary physical file is not also marked damaged. The following command backs up onsite data from the DISKDATA and DISKDIRS storage pool to the OFFDATA and OFFDIRS storage pools respectively. This task can also be automated using a server script as shown in 11.2.1, “Defining an Offsite Backup Schedule” on page 186.

```
adsm>backup stgpool diskdata offdata
```

```
ANR2017I Administrator SERVER_CONSOLE issued command: BACKUP STGPOOL diskdata  
offdata  
ANR0984I Process 5 for BACKUP STORAGE POOL started in the BACKGROUND at  
18:06:09.  
ANR2110I BACKUP STGPOOL started as process 5.  
ANR1210I Backup of primary storage pool DISKDATA to copy storage pool OFFDATA  
started as process 5.  
ANR8337I DLT volume V00003 mounted in drive DRV1 (/dev/mt1).  
ANR1340I Scratch volume V00003 is now defined in storage pool OFFDATA.  
ANR1212I Backup process 5 ended for storage pool DISKDATA.  
ANR0986I Process 5 for BACKUP STORAGE POOL running in the BACKGROUND processed  
16 items for a total of 35,639,296 bytes with a completion state of SUCCESS at  
18:08:07.  
ANR1214I Backup of primary storage pool DISKDATA to copy storage pool OFFDATA  
has ended. Files Backed Up: 16, Bytes Backed Up: 35639296, Unreadable Files:  
0, Unreadable Bytes: 0.
```

```
adsm> query volume * stg=offdata
```

Volume Name	Storage Pool Name	Device Class Name	Estimated Capacity (MB)	Pct Util	Volume Status
V00003	OFFDATA	CDLT	20,480.0	0.0	Filling

```
adsm>backup stgpool diskdirs offdirs
```

```
ANR2017I Administrator SERVER_CONSOLE issued command: BACKUP STGPOOL diskdirs  
offdirs  
ANR0984I Process 6 for BACKUP STORAGE POOL started in the BACKGROUND at  
07:34:49.  
ANR2110I BACKUP STGPOOL started as process 6.  
ANR1210I Backup of primary storage pool DISKDIRS to copy storage pool OFFDIRS  
started as process 6.  
ANR8337I DLT volume V00001 mounted in drive DRV1 (/dev/mt1).  
ANR1212I Backup process 6 ended for storage pool DISKDIRS.  
ANR0986I Process 6 for BACKUP STORAGE POOL running in the BACKGROUND processed  
24 items for a total of 507,904 bytes with a completion state of SUCCESS at  
07:36:23.  
ANR1214I Backup of primary storage pool DISKDIRS to copy storage pool OFFDIRS  
has ended. Files Backed Up: 24, Bytes Backed Up: 507904, Unreadable Files: 0,  
Unreadable Bytes: 0.
```

```
adsm> query volume * stg=offdirs
```

Volume Name	Storage Pool Name	Device Class Name	Estimated Capacity (MB)	Pct Util	Volume Status
V00001	OFFDIRS	CDLT	20,480.0	0.0	Filling

12.4.2.2 List Volumes to Send to Offsite

You must use the `query volume` command to list all offsite tapes that need to go to the offsite location. In our example, we assume that an offsite tape is any tape that belongs to the copy storage called OFFDATA and OFFDIRS, and thus, must be moved to a separate safe location (a vault).

```

adsm> query volume * access=readwrite,readonly status=full,filling
stgpool=offdata

Volume Name          Storage      Device      Estimated   Pct   Volume
                    Pool Name   Class Name  Capacity    Util  Status
                    -----
V00003                OFFDATA     CDLT        20,480.0    0.0  Filling

adsm> query volume * access=readwrite,readonly status=full,filling
stgpool=offdirs

Volume Name          Storage      Device      Estimated   Pct   Volume
                    Pool Name   Class Name  Capacity    Util  Status
                    -----
V00001                OFFDIRS     CDLT        20,480.0    0.0  Filling

```

As you can see from the example, the `query volume` listed two volumes (V00003 and V00001). The offsite movement operation must take into account two types of volumes: FULL and FILLING. You must issue the command for both status, so that you move all needed volumes.

We recommend that you keep a record of these volumes for tracking purposes.

12.4.2.3 Checking Out Volumes from the Library

This command is only used if you have an automated tape library. Use the `checkout libvolume` commands to remove ADSM control of a storage volume located in a library. You must use the same volume listing that you got from the previous `query volume` section. In our case, the volumes are V00003 and V00001:

```

adsm>checkout libvolume ldlt v00003
ANR2017I Administrator SERVER_CONSOLE issued command: CHECKOUT libv ldlt v00003

ANR0984I Process 8 for CHECKOUT LIBVOLUME started in the BACKGROUND at
07:42:16.
ANR8434I CHECKOUT LIBVOLUME: Operation for volume V00003 in library LDLT
started as process 8.
ANR8336I Verifying label of DLT volume V00003 in drive DRV1 (/dev/mt1).
ANR8307I 001: Remove DLT volume V00003 from slot with element number 3 of
library LDLT; issue 'REPLY' along with the request ID when ready.

adsm>reply 1
ANR2017I Administrator SERVER_CONSOLE issued command: REPLY 1
ANR8499I Command accepted.
ANR8438I CHECKOUT LIBVOLUME for volume V00003 in library LDLT completed
successfully.
ANR0985I Process 8 for CHECKOUT LIBVOLUME running in the BACKGROUND completed
with completion state SUCCESS at 07:44:08.

adsm>checkout libvolume ldlt v00001
ANR2017I Administrator SERVER_CONSOLE issued command: CHECKOUT libv ldlt v00001

ANR0984I Process 9 for CHECKOUT LIBVOLUME started in the BACKGROUND at
07:45:17.
ANR8434I CHECKOUT LIBVOLUME: Operation for volume V00001 in library LDLT
started as process 9.
ANR8336I Verifying label of DLT volume V00001 in drive DRV1 (/dev/mt1).
ANR8307I 002: Remove DLT volume V00001 from slot with element number 5 of
library LDLT; issue 'REPLY' along with the request ID when ready.

adsm>reply 2
ANR2017I Administrator SERVER_CONSOLE issued command: REPLY 2
ANR8499I Command accepted.
ANR8438I CHECKOUT LIBVOLUME for volume V00001 in library LDLT completed
successfully.
ANR0985I Process 9for CHECKOUT LIBVOLUME running in the BACKGROUND completed
with completion state SUCCESS at 07:48:09.

```

12.4.2.4 Update Volume Location to Offsite

Before you update the volume information to the offsite location, you have to make sure that the desired volumes have been physically sent and delivered to the offsite location. To avoid a situation in which the volumes are requested for mounts while in transition to the offsite location, you can set the volumes to unavailable while in transition, by using the `update volume` command with `STATUS` set to `UNAVAILABLE`.

You can change the offsite location name to one suitable for your site. Once in the offsite location, ADSM never asks for volumes with a status of offsite to be mounted.

The next example assumes that you have the volumes `V00003` and `V00001` from the previous `backup stgpool` commands and the volumes have been delivered to the offsite location. Now you want to update their location to "ITSO vault":

```
adsm> update volume * access=offsite location="ITSO vault" \  
cont> whereaccess=readwrite,readonly wherestg=offdata wherestatus=full,filling  
ANR2207I Volume V00003 updated.  
  
adsm> query volume v00003 format=detail  
  
          Volume Name: V00003  
          Storage Pool Name: OFFDATA  
          Device Class Name: CDLT  
          Estimated Capacity (MB): 20,480.0  
          Pct Util: 0.0  
          Volume Status: Filling  
          Access: Offsite  
          Pct. Reclaimable Space: 100.0  
          Scratch Volume?: No  
          In Error State?: No  
          Number of Writable Sides: 1  
          Number of Times Mounted: 1  
          Write Pass Number: 1  
          Approx. Date Last Written: 03/05/1999 07:14:51  
          Approx. Date Last Read: 03/05/1999 07:14:48  
          Date Became Pending:  
          Number of Write Errors: 0  
          Number of Read Errors: 0  
          Volume Location: ITSO vault  
          Last Update by (administrator): ARMANDO  
          Last Update Date/Time: 03/05/1999 08:09:50  
  
adsm> update volume * access=offsite location="ITSO vault" \  
cont> whereaccess=readwrite,readonly wherestg=offdirs wherestatus=full,filling  
ANR2207I Volume V00001 updated.  
  
adsm> query volume v00001 format=detail  
  
          Volume Name: V00001  
          Storage Pool Name: OFFDIRS  
          Device Class Name: CDLT  
          Estimated Capacity (MB): 20,480.0  
          Pct Util: 0.0  
          Volume Status: Filling  
          Access: Offsite  
          Pct. Reclaimable Space: 100.0  
          Scratch Volume?: No  
          In Error State?: No  
          Number of Writable Sides: 1  
          Number of Times Mounted: 1  
          Write Pass Number: 1  
          Approx. Date Last Written: 03/05/1999 07:36:22  
          Approx. Date Last Read: 03/05/1999 07:36:18  
          Date Became Pending:  
          Number of Write Errors: 0  
          Number of Read Errors: 0  
          Volume Location: ITSO vault  
          Last Update by (administrator): ARMANDO  
          Last Update Date/Time: 03/05/1999 08:11:42
```

12.4.3 Offsite Tape Management to Onsite

In this section we show how to move offsite volumes to the onsite location. These offsite volumes either have no active data on them (status of EMPTY) or are being returned to recreate damaged data in a primary storage pool.

Offsite tapes can be seen as your second good copy of critical data. As you perform backup and archive operations, all data is saved in storage pools, which are the basic unit for onsite operations. When you perform a `backup stgpool` command, AD SM verifies which data has not been sent to offsite yet and then performs the copy operation. Since this is an incremental storage duplication, only part of the whole onsite data is actually copied and thus, makes it time-effective and cost-effective (fewer tapes needed).

As backup and archive data age and expire, the offsite tapes contain less active data. Offsite tape reclamation consolidates these tape volumes, creating full tapes and empty tapes in the process. These tapes remain in a *pending* state until after the reuse delay period has expired. At this time the status of the tape changes to EMPTY. Empty tapes can be returned onsite.

The steps that you need to perform to receive tapes from offsite are the following:

1. Identify the volumes to be brought onsite.
2. Move the tape volumes from offsite to onsite.
3. Update the volume location.
4. Check in storage volumes into a library.

12.4.3.1 Listing Volumes to Bring Onsite

You must first check which volumes are available to return onsite. The following `query volume` command shows two volumes (V00003 and V00001) that are now available to be brought back:

```
adsm> query volume * access=offsite status=empty
```

Volume Name	Storage Pool Name	Device Class Name	Estimated Capacity (MB)	Pct Util	Volume Status
V00001	OFFDIRS	CDLT	0.0	0.0	Empty
V00003	OFFDATA	CDLT	0.0	0.0	Empty

We recommend that you keep a record of these volumes for tracking purposes and use it as a shipping list for your offsite location.

The tape volumes should now be moved from the offsite location to your onsite location.

12.4.3.2 Updating Volume Locations to Onsite

When the previously identified tape volumes have been returned onsite, you must update the location and access mode, so that the volume can be reused.

```

adsm> update volume * access=readwrite location="" wherestg=offdata \
cont> wherestatus=empty whereaccess=offsite
ANR2207I Volume V00003 updated.

adsm> query volume v00003 format=detail

          Volume Name: V00003
          Storage Pool Name: OFFDATA
          Device Class Name: CDLT
          Estimated Capacity (MB): 0.0
          Pct Util: 0.0
          Volume Status: Empty
          Access: Read/Write
          Pct. Reclaimable Space: 0.0
          Scratch Volume?: No
          In Error State?: No
          Number of Writable Sides: 1
          Number of Times Mounted: 1
          Write Pass Number: 1
          Approx. Date Last Written: 03/05/1999 07:14:51
          Approx. Date Last Read: 03/05/1999 07:14:48
          Date Became Pending:
          Number of Write Errors: 0
          Number of Read Errors: 0
          Volume Location:
          Last Update by (administrator): ARMANDO
          Last Update Date/Time: 03/05/1999 08:37:32

adsm> update volume * access=readwrite location="" wherestg=offdirs \
cont> wherestatus=empty whereaccess=offsite
ANR2207I Volume V00001 updated.

adsm> query volume v00001 format=detail

          Volume Name: V00001
          Storage Pool Name: OFFDIRS
          Device Class Name: CDLT
          Estimated Capacity (MB): 0.0
          Pct Util: 0.0
          Volume Status: Empty
          Access: Read/Write
          Pct. Reclaimable Space: 0.0
          Scratch Volume?: No
          In Error State?: No
          Number of Writable Sides: 1
          Number of Times Mounted: 1
          Write Pass Number: 1
          Approx. Date Last Written: 03/05/1999 07:36:22
          Approx. Date Last Read: 03/05/1999 07:36:18
          Date Became Pending:
          Number of Write Errors: 0
          Number of Read Errors: 0
          Volume Location:
          Last Update by (administrator): ARMANDO
          Last Update Date/Time: 03/05/1999 08:52:51

```

12.4.3.3 Checking in Storage Volumes into a Library

This command may only be used if you have an automated tape library. Use the `checkin libvolume` command to add a sequential access storage volume to ADSM's inventory for an automated library. This command informs the ADSM server that a volume is available for use. The server does not use any volumes that physically reside in an automated library until a volume has been checked in via the `checkin libvolume` command. The volumes are checked in as scratch because they are empty.

```
adsm>checkin libvolume ldlt search=yes status=scratch  
  
ANR2017I Administrator SERVER_CONSOLE issued command: CHECKIN libv ldlt  
search=yes status=scratch  
ANR0984I Process 3 for CHECKIN LIBVOLUME started in the BACKGROUND at 17:55:16.  
ANR8422I CHECKIN LIBVOLUME: Operation for library LDLT started as process 3.  
ANR8430I Volume V00001 has been checked into library LDLT.  
ANR8430I Volume V00003 has been checked into library LDLT.  
ANR8431I CHECKIN LIBVOLUME process completed for library LDLT; 2 volume(s)  
found.  
ANR0985I Process 3 for CHECKIN LIBVOLUME running in the BACKGROUND completed  
with completion state SUCCESS at 17:59:42.
```

12.4.4 Reclaiming Offsite Tapes

We can think of a moment in time where onsite data gets expired, and thus, must reflect the same situation offsite. ADSM does not need the physical offsite tape to perform this operation. Because ADSM has its own database (which is actually the index for all files), it is just a matter of updating the new data as being "deleted" in the database.

After some time, you may need to run a reclaim procedure, so that unused space is released. Onsite tapes are merged so that you can release tapes for use. Offsite tapes are a little bit different due to their own nature. To recycle an offsite tape, you need first to "create" a new good onsite volume copy of all remaining data. This is our new consolidated offsite volume. After that, we need to send this new tape before actually reclaiming the old offsite volumes. ADSM handles this situation by creating one or more tapes from the onsite ones.

ADSM will only ask to bring the old (unused offsite) tape after you send the new (and now true) tape to offsite.

12.4.5 Database Backup Management

The management of ADSM database backups is a critical task in the management of your environment. Backups of the database should be done every day, and the copies should be moved offsite to provide disaster recovery.

The steps involved in the movement of backups from onsite to offsite, and their return, are similar to those for storage pool data. The Disaster Recovery Manager (DRM) feature simplifies the management of onsite-offsite tape movement for both storage pool and database data. We highly recommend using DRM.

12.4.5.1 Moving Database Backups Offsite

The steps that you need to perform to send database backups offsite are:

1. Back up your database.
2. Identify the volumes to be taken offsite.
3. Remove volumes from the library.
4. Move the volumes offsite.

Back Up Your Database

You can perform a backup of the database at any time, by using the `backup db` command. For examples, refer to 4.8, “Database Backup” on page 62.

This task can also be automated using a server script as shown in 11.2.1, “Defining an Offsite Backup Schedule” on page 186.

Database backup volumes are tracked in a different manner to those of storage pools. You can use the `query volhistory` command to identify tape volumes containing database backups. Alternately, you can use the activity log messages to identify those tapes. Assuming that the database backup has been done, the command to display the volumes is:

```
adsm>query volhistory type=dbbackup begindate=today
      Date/Time: 03/10/99 10:48:14
      Volume Type: BACKUPFULL
      Backup Series: 13
      Backup Operation: 0
      Volume Seq: 1
      Device Class: CDLT
      Volume Name: DB0003
      Volume Location:
```

The volumes you require are those identified in the database backup with the latest time stamp.

Remove Backups from Library

This step only applies if your backup tapes are in a library.

You can use the `checkout libvolume` command to remove the database backup tapes from the library. This is the same process as that for storage pool data.

```
adsm>checkout libvolume ldlt db0003
ANR2017I Administrator PHIL issued command: CHECKOUT libv ldlt db0003

ANR0984I Process 8 for CHECKOUT LIBVOLUME started in the BACKGROUND at
07:55:16.
ANR8434I CHECKOUT LIBVOLUME: Operation for volume DB0003 in library LDLT
started as process 8.
ANR8336I Verifying label of DLT volume DB0003 in drive DRV1 (/dev/mt1).
ANR8438I CHECKOUT LIBVOLUME for volume DB0003 in library LDLT completed
successfully.
ANR0985I Process 8 for CHECKOUT LIBVOLUME running in the BACKGROUND completed
with completion state SUCCESS at 07:58:05.
```

Move Volumes Offsite

You should move the backup volumes to your offsite location. We recommend that you track the tapes in a movement log.

12.4.5.2 Reusing Backup Tapes

You should keep a minimum of three days of database backup tapes. If you want to allow for holiday weekends, then a larger number such as four or five may be more appropriate. Our redbook environment uses five tapes.

The steps that you need to perform to reuse database backup tapes are:

1. Identify the volumes to be brought onsite.
2. Move the tape volumes from offsite to onsite.

3. Update the volume to scratch.
4. Check in storage volumes into a library.

Identify the Volumes to be Brought Onsite

Database backups remain active until their entry is deleted from the volume history. However, once it has been deleted, ADSM has no knowledge of the volume at all.

As there is no way to display just the oldest database backup volumes, we use the `query volhistory` command to identify all tape volumes containing database backups older than a specific number of days. The volumes you require are those with the oldest time stamp. The following example uses five days.

```
adsm> query volhistory type=dbbackup enddate=today-5
      Date/Time: 03/05/1999 06:27:10
      Volume Type: BACKUPFULL
      Backup Series: 17
      Backup Operation: 0
      Volume Seq: 1
      Device Class: CDLT
      Volume Name: DB0001
      Volume Location:
```

Move the Tape Volumes from Offsite to Onsite

You should move the backup volumes from your offsite location to your site. We recommend that you track the tapes in a movement log.

Update the Volume to Scratch

You use the `delete volhistory` command to return the database backup volumes to scratch. As there is no way to just delete the oldest database backup volumes, you must delete backups older than a number of days. The following example uses five days.

```
adsm> delete volhistory type=dbbackup todate=today-5
Do you wish to proceed? (Yes/No) y
ANR2467I DELETE VOLHISTORY: 0 sequential volume history entries were
successfully deleted.
```

Check in Storage Volumes into a Library

This step only applies if your backup tapes are in a library.

You can use the `checkin libvolume` command to put the database backup tapes into the library. This is the same process as that for storage pool data.

```
adsm>checkin libvolume ldlt search=yes status=scratch
ANR2017I Administrator AMANDA issued command: CHECKIN libv ldlt
search=yes status=scratch
ANR0984I Process 3 for CHECKIN LIBVOLUME started in the BACKGROUND at 19:59:55.
ANR8422I CHECKIN LIBVOLUME: Operation for library LDLT started as process 3.
ANR8430I Volume DB0003 has been checked into library LDLT.
ANR8431I CHECKIN LIBVOLUME process completed for library LDLT; 1 volume(s)
found.
ANR0985I Process 3 for CHECKIN LIBVOLUME running in the BACKGROUND completed
with completion state SUCCESS at 20:01:42.
```

12.5 Error Conditions

ADSM can write error information to both the ADSM activity log or the operating system log. Some errors may not be related to ADSM. Thus, it may detect a system failure which may be related to hardware problems, or even a software component reporting error conditions.

12.5.1 ADSM Errors

Error analysis is a key point to your ADSM environment. If for any reason ADSM is not able to write data to a tape, it may retry the operation (using hardware built-in recovery features) or it may fail due to a non-recoverable error. ADSM will notify all error conditions to the console, the ADSM activity log, or the operating system log. For example, you can check that ADSM has an error condition logged by using the `query actlog` command with the `SEARCH` parameter:

```
adsm>query actlog begindate=02/09/1999 search=*error*

ANR2017I Administrator SERVER_CONSOLE issued command: QUERY ACTLOG begindate=0-
2/09/1999 search=*error*

Date/Time          Message
-----
...
02/10/99 11:06:01  ANR8300E I/O error on library LDLT (OP=00006C03, CC=207,
KEY=05, ASC=21, ASCQ=01, SENSE=70.00.05.00.00.00.0E.0-
0.00.00.00.21.01.00.00.00.00.38.00., Description=Device
is not in a state capable of performing request). Refer
to Appendix B in the 'Messages' manual for recommended
action.
02/10/99 11:06:02  ANR8300E I/O error on library LDLT (OP=00006C03, CC=314,
KEY=05, ASC=3B, ASCQ=0E, SENSE=70.00.05.00.00.00.0E.0-
0.00.00.00.3B.0E.00.00.00.00.38.00., Description=The
source slot or drive was empty in an attempt to move a
volume). Refer to Appendix B in the 'Messages' manual
for recommended action.
...
02/11/99 00:17:09  ANR8353E 004: I/O error reading label of volume in drive
DRV0 (/dev/mt0).
...
```

As you can see from the previous example, there are two library errors and one volume error. The volume error may be because of a true hardware problem or a media error. In this case, you may try to use another tape to isolate the problem or even use another tape unit. You might realize that it was just a matter of cleaning the unit. In our example, ADSM is not communicating with the library autochanger when we perform a `REPLY` operation.

12.5.2 Machine Errors

Although ADSM can handle many error conditions, it cannot resolve all possible errors. Therefore, when an unexpected error condition is met, ADSM notifies the operating system event logging so that you may further isolate the problem. Note that not all error messages are true problems. Sometimes, they are simply information messages (for example, if a 3590 tape needs cleaning, then the 3494 library manager may mount the tape and create an information record indicating that a `CLEAN` operation was performed).

12.5.2.1 AIX errpt

You can locate ADSM errors on the AIX error log by using the operating system command `errpt`. You also can use the system management interface tool (SMIT) fast path to run the `errpt` command. To use the SMIT fast path, enter SMIT `ERRPT`.

Note that the information shown in the `errpt` may not be all from ADSM. In this example, the `LFTDD` error is not related to ADSM. The `lb0` and `mt` errors may be an ADSM problem or a hardware problem. You and/or the technical staff must investigate the conditions when the problem occurs and take the required actions.

```
kindu:[/]$ errpt | pg
IDENTIFIER  TIMESTAMP  T C RESOURCE_NAME  DESCRIPTION
4225DB66    0205183999 T H lb0           OPERATOR INTERVENTION REQUIRED
4225DB66    0205175599 T H lb0           OPERATOR INTERVENTION REQUIRED
4225DB66    0205175499 T H lb0           OPERATOR INTERVENTION REQUIRED
4225DB66    0204175899 T H lb0           OPERATOR INTERVENTION REQUIRED
4225DB66    0204175299 T H lb0           OPERATOR INTERVENTION REQUIRED
E85C5C4C    0204141899 P S LFTDD          SOFTWARE PROGRAM ERROR
2BFA76F6    0204141499 T S SYSPROC       SYSTEM SHUTDOWN BY USER
9DFCFDEE    0204141799 T O errdemon      ERROR LOGGING TURNED ON
192AC071    0204141399 T O errdemon      ERROR LOGGING TURNED OFF
4225DB66    0203152899 T H mt0           OPERATOR INTERVENTION REQUIRED
4225DB66    0203145899 T H lb0           OPERATOR INTERVENTION REQUIRED
4225DB66    0203144799 T H mt1           OPERATOR INTERVENTION REQUIRED
```

You can get additional details by using the `-a` flag:

```
kindu:[/]$ errpt -a
...
-----
LABEL:          ADSM_DD_LOG4
IDENTIFIER:     4225DB66

Date/Time:      Fri Feb  5 18:39:56
Sequence Number: 355
Machine Id:     000034217600
Node Id:        kindu
Class:          H
Type:           TEMP
Resource Name:  lb0
Resource Class: library
Resource Type:  ADSM-SCSI-LB
Location:       00-00-0S-3,0

Description
OPERATOR INTERVENTION REQUIRED

Probable Causes
ATTACHED SCSI TARGET DEVICE

Failure Causes
ATTACHED SCSI TARGET DEVICE
:

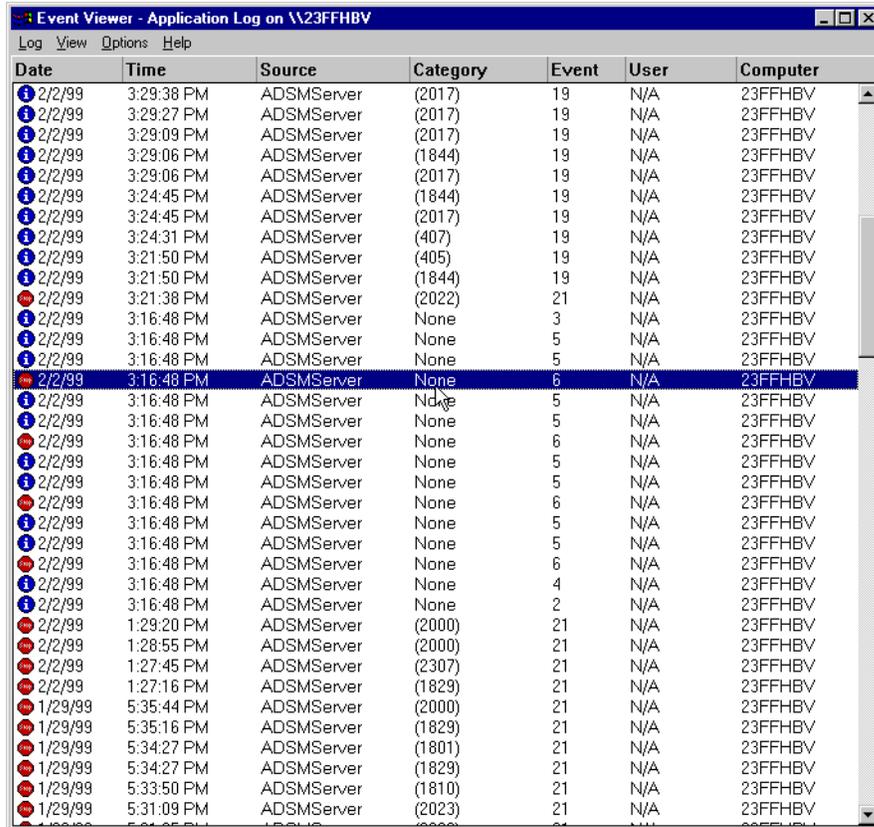
Recommended Actions
CORRECT THEN RETRY
CHECK POWER
CHECK PHYSICAL INSTALLATION

...

```

12.5.2.2 Windows NT Event Viewer

ADSM logs error, information, and warning conditions to the Windows NT application log. You can browse it by using the Event Viewer utility which comes with Windows NT. Figure 39 on page 228 shows you an example of a Windows NT error screen with some ADSM error. Note that not all of them may be true error conditions. You must evaluate if it is actually a problem or simply a misconfiguration.



The screenshot shows the Windows Event Viewer window titled "Event Viewer - Application Log on \\23FFHBY". The window displays a list of events in a table format. The columns are Date, Time, Source, Category, Event, User, and Computer. The events are sorted by time, with the most recent at the top. A red icon in the Date column indicates an error event. The selected event is dated 2/2/99 at 3:16:48 PM, with Source ADSMServer, Category None, Event 6, User N/A, and Computer 23FFHBY.

Date	Time	Source	Category	Event	User	Computer
2/2/99	3:29:38 PM	ADSMServer	(2017)	19	N/A	23FFHBY
2/2/99	3:29:27 PM	ADSMServer	(2017)	19	N/A	23FFHBY
2/2/99	3:29:09 PM	ADSMServer	(2017)	19	N/A	23FFHBY
2/2/99	3:29:06 PM	ADSMServer	(1844)	19	N/A	23FFHBY
2/2/99	3:29:06 PM	ADSMServer	(2017)	19	N/A	23FFHBY
2/2/99	3:24:45 PM	ADSMServer	(1844)	19	N/A	23FFHBY
2/2/99	3:24:45 PM	ADSMServer	(2017)	19	N/A	23FFHBY
2/2/99	3:24:31 PM	ADSMServer	(407)	19	N/A	23FFHBY
2/2/99	3:21:50 PM	ADSMServer	(405)	19	N/A	23FFHBY
2/2/99	3:21:50 PM	ADSMServer	(1844)	19	N/A	23FFHBY
2/2/99	3:21:38 PM	ADSMServer	(2022)	21	N/A	23FFHBY
2/2/99	3:16:48 PM	ADSMServer	None	3	N/A	23FFHBY
2/2/99	3:16:48 PM	ADSMServer	None	5	N/A	23FFHBY
2/2/99	3:16:48 PM	ADSMServer	None	5	N/A	23FFHBY
2/2/99	3:16:48 PM	ADSMServer	None	6	N/A	23FFHBY
2/2/99	3:16:48 PM	ADSMServer	None	5	N/A	23FFHBY
2/2/99	3:16:48 PM	ADSMServer	None	5	N/A	23FFHBY
2/2/99	3:16:48 PM	ADSMServer	None	6	N/A	23FFHBY
2/2/99	3:16:48 PM	ADSMServer	None	5	N/A	23FFHBY
2/2/99	3:16:48 PM	ADSMServer	None	6	N/A	23FFHBY
2/2/99	3:16:48 PM	ADSMServer	None	4	N/A	23FFHBY
2/2/99	3:16:48 PM	ADSMServer	None	2	N/A	23FFHBY
2/2/99	1:29:20 PM	ADSMServer	(2000)	21	N/A	23FFHBY
2/2/99	1:28:55 PM	ADSMServer	(2000)	21	N/A	23FFHBY
2/2/99	1:27:45 PM	ADSMServer	(2307)	21	N/A	23FFHBY
2/2/99	1:27:16 PM	ADSMServer	(1829)	21	N/A	23FFHBY
1/29/99	5:35:44 PM	ADSMServer	(2000)	21	N/A	23FFHBY
1/29/99	5:35:16 PM	ADSMServer	(1829)	21	N/A	23FFHBY
1/29/99	5:34:27 PM	ADSMServer	(1801)	21	N/A	23FFHBY
1/29/99	5:34:27 PM	ADSMServer	(1829)	21	N/A	23FFHBY
1/29/99	5:33:50 PM	ADSMServer	(1810)	21	N/A	23FFHBY
1/29/99	5:31:09 PM	ADSMServer	(2023)	21	N/A	23FFHBY

Figure 39. ADSM Error Entries in Windows Event Viewer

In this example we selected one of the red events, which is an error condition. Figure 40 on page 229 shows us that a severe error occurred (most probably, an attempt to start ADSM when it was already running).

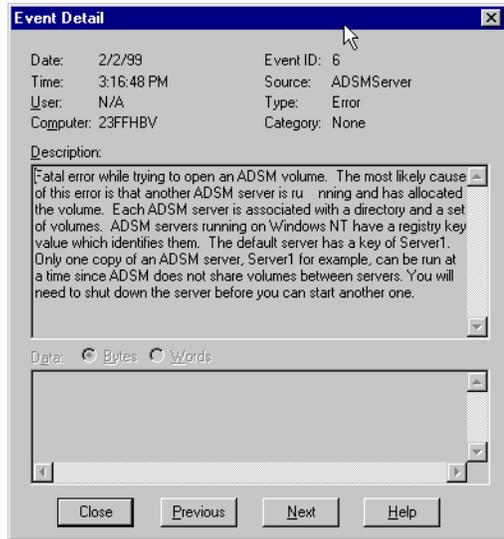


Figure 40. Detailed Event Information of ADSM Error

Chapter 13. Performance Considerations

This chapter focuses on some parameters that should be tuned to get the maximum performance during backup, archive, restore and retrieve operations. The first section is a discussion of the parameters applicable on the ADSM server. The next section covers the client node parameters.

More detailed explanations about performance tuning can be found in the ADSM Version 3 Release 1 Performance Tuning Guide at <http://books.adsm.org>.

13.1 ADSM Server

This section is a discussion of the tuning parameters applicable to all ADSM servers. There are other tuning parameters that can also be considered, but due to the scope limitations of this guide, we will not be able to discuss all of them.

Note:

The following parameters must be set in the server options file. You need to restart the ADSM server for the changes to take effect.

13.1.1 BUFPOOLSIZE

Cache storage is provided by the database buffer poolsize, which allows the database pages to remain in memory for longer periods of time. This will entail the database pages to remain in cache. The server can make continuous updates to the pages without requiring I/O operations to external storage. While a database buffer pool can improve server performance, it will also require more virtual memory.

An optimal setting for the database buffer pool is one in which the cache hit percentage is greater than or equal to 98%. To check the cache hit percentage, use the `query db` command with `FORMAT=DETAILED`.

```
adsm> query db format=detail

    Available Space (MB): 40
    Assigned Capacity (MB): 40
    Maximum Extension (MB): 0
    Maximum Reduction (MB): 32
      Page Size (bytes): 4,096
    Total Usable Pages: 10,240
      Used Pages: 1,636
        Pct Util: 16.0
      Max. Pct Util: 16.0
    Physical Volumes: 12
    Buffer Pool Pages: 512
    Total Buffer Requests: 46,315
      Cache Hit Pct.: 99.67
      Cache Wait Pct.: 0.00
    Backup in Progress?: No
    Type of Backup In Progress:
    Incrementals Since Last Full: 0
    Changed Since Last Backup (MB): 0.25
      Percentage Changed: 3.91
    Last Complete Backup Date/Time: 02/13/1999 00:23:46
```

Increasing the BUFPOOLSIZE parameter can improve the performance of many ADSM server functions such as multi-client backup, storage pool migration, storage pool backup, expiration processing, and move data. If the cache hit percentage is lower than 99%, increase the size of the BUFPOOLSIZE parameter in the server options file. For most servers, we recommend starting with a value of 32768, which equals 8192 database pages. If you have enough memory, increase in 1MB increments. A cache hit percentage greater than 99% is an indication that the proper BUFPOOLSIZE has been reached. However, continuing to raise BUFPOOLSIZE beyond that level can be very helpful. While increasing BUFPOOLSIZE, care must be taken not to cause paging in the virtual memory system. Monitor system memory usage to check for any increased paging after the BUFPOOLSIZE change. The recommended values for this parameter are shown in Table 43 on page 232.

Table 43. Recommended Values for the BUFPOOLSIZE

System memory	Recommended BUFPOOLSIZE (KB)
32	2048
48	3072
64	4096
96	9216
128	14336
160	20480
256	32768
512	65536
1024	131072

Use the `reset bufpool` command to reset the cache hit statistics.

```
adsm> reset bufpool
ANR0381I Bufferpool statistics were successfully reset.
```

13.1.2 EXPINTERVAL

The ADSM server runs automatic inventory expiration, and this option specifies the interval in hours for this process. Inventory expiration removes client backup and archive file copies from the server.

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 disk storage space.

Expiration processing is very CPU intensive. If possible, it should be run when other ADSM processes are not occurring. To enable this, either schedule expiration once per day, or set EXPINTERVAL to 0 and manually start the process with the `expire inventory` command at the server.

13.1.3 LOGPOOLSIZE

This parameter specifies the size of the recovery log buffer pool size in kilobytes. A large recovery log buffer pool may increase the rate by which recovery log transactions are committed to the database, but it also requires more memory. 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 the frequency in which the server forces records to the recovery log. To determine if LOGPOOLSIZE should be increased, monitor the value of Log Pool Percentage Wait. To check the wait percentage, use the `query log` command with `FORMAT=DETAIL`.

```
adsm> query log format=detail

    Available Space (MB): 108
    Assigned Capacity (MB): 24
    Maximum Extension (MB): 84
    Maximum Reduction (MB): 16
    Page Size (bytes): 4,096
    Total Usable Pages: 5,632
    Used Pages: 351
    Pct Util: 6.2
    Max. Pct Util: 7.3
    Physical Volumes: 2
    Log Pool Pages: 128
    Log Pool Pct. Util: 0.23
    Log Pool Pct. Wait: 0.00
    Cumulative Consumption (MB): 493.24
    Consumption Reset Date/Time: 02/13/1998 11:00:01
```

If the value is greater than zero, increase the value of LOGPOOLSIZE. As the size of the recovery log buffer pool is increased, remember to monitor system memory usage.

13.1.4 MAXSESSION

The MAXSESSION parameter specifies the maximum number of simultaneous client sessions that can connect with the ADSM server. The default value is 25 client sessions. The minimum value is 2 client sessions. The maximum value is limited only by available virtual memory or communication resources. This parameter specifies the maximum number of simultaneous client sessions that can connect with the ADSM Server. By limiting the number of clients, server performance can be improved, but the availability of ADSM services to the clients is reduced.

13.1.5 MOVEBATCHSIZE and MOVESIZETHRESH

These two options, MOVEBATCHSIZE and MOVESIZETHRESH, help tune the performance of the server processes that involve the movement of data between storage media. These processes include storage pool backup and restore, migration, reclamation, and move data.

These options specify the number of files that are to be moved and grouped together in a batch, within the same server transaction.

The number of client files moved for each server database transaction during a server storage pool backup or restore, migration, reclamation, or move data operation will be determined by the number and size of the files in the batch. If the number of files in the batch equals the MOVEBATCHSIZE before the

cumulative size of the files becomes greater than the MOVESIZETHRESH, then the MOVEBATCHSIZE is used to determine the number of files moved or copied in the transaction. If the cumulative size of files being gathered for a move or copy operation exceeds the MOVESIZETHRESH value before the number of files becomes equivalent to the MOVEBATCHSIZE, then the MOVESIZETHRESH value is used to determine the number of files moved or copied in the transaction.

Note that when the MOVEBATCHSIZE or MOVESIZETHRESH parameters are increased from their default values, the server will require more space in the recovery log. The recovery log may require an allocation space 2 or more times larger than a recovery log size which uses the defaults. In addition, the server requires a longer initialization time at startup. The impact of a larger recovery log size will be felt while running the server with the logmode set to NORMAL (the default value). If you choose to increase these values for performance reasons, be sure to monitor recovery log usage during the first few storage pool backup/restore, migration, reclamation, or move data executions to ensure sufficient recovery log space is available.

13.1.6 USELARGEBUFFER

This parameter increases communication and device I/O buffers. Both the client-server communication buffer and disk device I/O buffers have been increased from 32 KB to 256 KB. The communication I/O buffer is used during data transfer with a client session, such as a backup session. The disk I/O buffer is used when data is read from or written to a disk storage pool.

Significant improvement in data transfer operations and CPU usage has been observed when the USELARGEBUFFER feature is enabled. Increasing the buffer sizes allows client-server communications and disk I/O to be more efficiently performed than environments without this feature enabled. Reads and writes are quicker and server resources are better utilized. By reducing CPU utilization, more clients can concurrently be serviced, improving overall system performance. With ADSM Version 3 these benefits are complemented with server aggregation. Aggregation groups smaller, logical client files into fewer but larger physical files at the server level. Larger files are better able to take advantage of the larger buffers.

The USELARGEBUFFER option is enabled by default and will result in the server storing data in a new format.

13.1.7 TCPWINDOWSIZE

This option specifies the size of the TCP sliding window in kilobytes. The TCPWINDOWSIZE option overrides the operating system's TCP send and receive spaces. In AIX for instance, these parameters are named tcp_sendspace and tcp_recvspace. The TCPWINDOWSIZE option specifies the size of the TCP sliding window for all clients and all but MVS servers. A larger window size can improve communication performance, but uses more memory. It enables multiple frames to be sent before an acknowledgment is obtained from the receiver. If long transmission delays are being observed, increasing the TCPWINDOWSIZE may improve throughput.

The size of the TCP/IP buffer is used when sending or receiving data. The window size used in a session is the smaller of the server and client window sizes. Larger window sizes use additional memory but may improve performance.

13.1.8 TXNGROUPMAX

The TXNGROUPMAX option specifies the number of files transferred as a group between commit points. This parameter is used in conjunction with the TXNBYTELIMIT client option. This option reduces the number of server transactions by increasing the number of files within any one transaction. Therefore, the amount of overhead during backup or restore caused by database commits is reduced.

When setting the size of transactions, consider setting a smaller size if you are suffering many resends due to files changing during backup when using static, shared static, or shared dynamic. This would apply to static as well as to shared because when the client realizes a file has changed during a backup and decides to not send it (the file, that is), the client would still have to resend the other files in that transaction.

13.2 Client Node

This section focuses on some tuning parameters on the ADSM client to obtain maximum performance when using ADSM. These parameters can be changed or added to the client options file.

13.2.1 COMPRESSION

The compression option compresses files before you send them to the ADSM server. Compressing your files decreases the amount of data storage that is required to store backup versions and archive copies of your files. It can, however, affect ADSM throughput.

Client data compression will save storage space, network capacity and server cycles. However, compression may have an adverse effect on throughput. If compression is on, throughputs can be significantly slower, depending on client processor speed, than if it were off. Compression may be beneficial for a fast processor on a slow network, but it is not for a slow processor on a fast network.

Files continue compression even if the file size continues to increase. To prevent continued compression if the filesize grows, and to send the file again without compression, use the COMPRESSALWAYS option set to NO.

This option controls compression only if your ADSM administrator specifies that your client node determines the selection using the `update node` command with the parameter COMPRESSION set to YES.

Two alternatives exist to using ADSM compression:

- If you are backing up to tape, and the tape drive supports its own compaction, use the tape drive compaction.
- Do not use ADSM compression if a client currently has built-in file compression support. ADSM compression on these clients will not yield additional reduction in the amount of data backed up to the server. The following platforms have built-in file compression: NetWare 4.x, Windows NT, and Windows 95.

13.2.2 COMPRESSALWAYS

The COMPRESSALWAYS option controls what occurs when a file grows during compression. You can continue compressing, or send the object again if it grows during compression. This option is used with the COMPRESSION option.

13.2.3 DIRMC

This option will have a big impact during restore. The DIRMC option specifies the management class you want ADSM to use for directories. The main benefit of managing directories separately from their files is faster file restore throughput from tape. The directories should be stored in a disk storage pool. The disk storage pool should not require a lot of space, since directories are typically very small.

In a restore operation, ADSM first restores the directories and then the files. This order of events can have a negative impact on restore from tape. For example, if you have slow tape storage pools, and the directories are stored on tape with the files, a large amount of time can be spent searching and mounting tapes to restore the directories. This time is saved with the DIRMC option by storing the directories in a disk storage pool.

13.2.4 QUIET

The QUIET option keeps messages from being written to the screen during ADSM backups. By default, ADSM displays information about each file it backs up. To prevent this, use the QUIET option. However, messages and summary information are still written to the log files.

Two main benefits exist in using the QUIET option:

- For tape backup, the first transaction group of data is always resent. To avoid this, use the QUIET option to reduce retransmissions at the client.
- If you are using the client scheduler to schedule ADSM backups, using the QUIET option dramatically reduces disk I/O overhead to the schedule log and improves ADSM throughput.

13.2.5 LARGECOMMBUFFERS

The LARGECOMMBUFFERS option specifies whether the client will use increased buffers to transfer large amounts of data between the client and the server. You can disable this option when your machine is running low on memory.

13.2.6 TCPBUFFSIZE

The TCPBUFFSIZE option specifies the size of the internal TCP communication buffer that is used to transfer data between the client node and the server. A large buffer can improve communication performance, but requires more memory.

13.2.7 TCPWINDOWSIZE

This option specifies the size of the TCP/IP sliding window in kilobytes. The TCPWINDOWSIZE option overrides the operating system's TCP send and receive spaces.

The TCPWINDOWSIZE option specifies the size of the TCP sliding window for all clients and all but MVS servers. A larger window size can improve communication performance, but uses more memory. It enables multiple frames to be sent before an acknowledgment is obtained from the receiver. If long transmission delays are being observed, increasing the TCPWINDOWSIZE may improve throughput.

13.2.8 TXNBYTELIMIT

Specifies the batch size, in kilobytes, for ADSM server transactions. The TXNBYTELIMIT option is used in conjunction with TXNGROUPMAX server option. This option reduces the number of server transactions by increasing the amount of data within any one transaction. Therefore, the amount of overhead during backup, restore, archive, and retrieve caused by database commits is reduced.

There are several items to consider when setting this parameter:

- Increasing the amount of data per transaction will increase recovery log requirements on the server. Check log and log pool space to ensure there is enough space. Also note that a larger log may result in longer server start-up times.
- Increasing the amount of data per transaction may result in more data being retransmitted if a retry occurs. This may negatively affect performance.
- The benefits of changing this parameter are dependent on configuration and workload characteristics. In particular, this parameter benefits tape storage pool backup more so than disk storage pool backup, especially if many small files are in the workload.

When setting the size of transactions, consider setting a smaller size if you are suffering many resends due to files changing during backup when using static, shared static, or shared dynamic. This would apply to static as well as to shared because when the client realizes a file has changed during a backup and decides to not send it, the file that is, it would still have to re-send the other files in that transaction.

13.3 How to Measure Performance

In measuring the performance of the ADSM solution, there is a need for the above mentioned parameters to be fine tuned and the system to be restarted. Fine tuning ADSM does not end with the backup, as the restore window is as important as the backup window. But in some cases where the backup parameters are needed in restoring the backed up data, then either the backup window or the restore window will have to give way a little.

In most cases where the backup parameters are different from the restore parameters, the backup and restore windows can be maximized.

To get the maximum performance of the system, a series of tests must be performed and recorded. It is only in this manner that the throughput can be tabulated. Tests should be conducted both for the backup and restore.

ADSM performance can be influenced by various tuning parameters. Tuning these functions for good performance requires diligence and expertise on the part of the installer. The number of parameters that may be set within ADSM are quite small; it is the tuning of the client, server, and network options that can become very complex.

Performance tuning for a single platform function is quite complex, but due to years of experience, it is generally well understood. ADSM, however, functions in the client-server realm, supports many operating systems, works across networks, and accepts different communication protocols. Consequently, there are many more factors that affect performance. The factors below can affect ADSM performance significantly:

- Client type
- Client speed
- Client activity
- Communication protocol type
- Communication protocol tuning
- Communication controller hardware
- Network activity
- Network speed
- Network reliability
- Server type
- Server speed
- Server activity
- Size and number of files
- Final output repository type (disk, tape, optical)

Clearly, with this many combinations, it is not feasible to discuss all possible combinations of these parameters within the scope of this guide.

Appendix A. Planning and Sizing Worksheets

The following is a collection of worksheets we introduce and recommend to use in Chapter 1, "ADSM Implementation Checklists" on page 1.

The redbook support material is available in softcopy on the internet from the redbooks Web server. Point your Web browser to:

<ftp://www.redbooks.ibm.com/redbooks/SG245416>

Alternatively you can go to:

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

and select Additional Redbook Materials (or follow the instructions given since the Web pages change frequently!).

Table 44. Client Requirements Worksheet

	Client 1	Client 2	Client 3	Client 4
Client name				
Contact information				
Operating system				
Total storage available (GB)				
Total storage used (GB)				
GB changed per backup				
Number of files backed up				
Data compression				
Backup window times				
Backup number of hours				
Required recovery time frame				
ADSM restore time frame				
GB copied per archive				
Number of files archived				
Number of archives kept				
Archive frequency				
Archive window times				
Archive number of hours				
Policy domain				
Client option set				

Table 45. Storage Policy Requirements Worksheet

	Example 1	Example 2	Example 3	Example 4
Group name				
Number of backup versions				
Backup file retention period				
Number of deleted versions				
Deleted file retention period				
Offsite copies				
Onsite collocation				
Offsite collocation				
Archive retention period				

Table 46. Database Worksheet

Filename (Primary)	Volume	Size (MB)	Filename (Copy)	Volume	Size (MB)
Total			Total		

Table 47. Recovery Log Worksheet

Filename (Primary)	Volume	Size (MB)	Filename (Copy)	Volume	Size (MB)
Total			Total		

Table 48. Primary Storage Pool Worksheet

Filename	Volume	Size (MB)
Total		

Table 49. Device Configuration and Volume History Worksheet

Name	Volume	Size (MB)
Total		

Table 50. Total ADSM Disk Required Worksheet

	Size (MB)
ADSM code (dependent on platform)	
ADSM database	
ADSM recovery log	
Primary storage pools	
Device configuration table and volume history table	
Other (RAID, Operating system)	
Total	

Table 51. Tape Drive Configuration Worksheet

	Option
Library model	
Number of drives	
Drive model	
Number of onsite tape volumes	
Number of offsite tape volumes	
Number of database backup volumes	
Number of scratch tapes	
Total tape volumes required	

Table 52. Administrator IDs Worksheet

Functions	ADSM ID	Authority
Server console	SERVER_CONSOLE	System

Table 53. License Requirement Worksheet

	Required	Required	Required
Server type			
Network communications			
Hierarchal storage manager			
Backup-Archive clients			
Advanced device support			
Enterprise administration			
Disaster recovery manager			
OpenSystems environment			
Server-to-server virtual volumes			

Appendix B. Redbook Support Material

This appendix contains listings of support material for this redbook. See section 1.2, "Redbook Support Material" on page 2 for considerations regarding this material.

The redbook support material is available in softcopy on the internet from the redbooks Web server. Point your Web browser to:

<ftp://www.redbooks.ibm.com/redbooks/SG245416>

Alternatively you can go to:

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

and select Additional Redbook Materials (or follow the instructions given since the Web pages change frequently!).

B.1 Macros

We have provided macros to help you implement your ADSM environment. The names and values within those macros reflect the recommendations for our redbook environment. You may change them to suit your environment as required.

B.1.1 Define Administrators

The following is an ADSM macro which contains the administrative commands to define our administrators. The name of the macro is `mac.admins`.

```
/*=====*/
/* ADSM V3 Redbook Macro - Define Administrators */
/* Getting Started with ADSM - Implementation (SG24-5416) */
/* */
/* These scripts are supplied to help you implement an ADSM */
/* environment. The names and values reflect the recommendations */
/* for our Redbook environment. You may change them to suit your */
/* environment as required. */
/* */
/* This file is designed to be run as an Administrative Command */
/* Line macro. You should use the -itemcommit parameter when you */
/* start the admin session, otherwise the macro could fail!! */
/* */
/*=====*/

/*-----*/
/* System Administrators */
/*-----*/
register admin sysadmin sysadmin contact='System Administrator'
grant authority sysadmin classes=system

/*-----*/
/* System Support Administrators */
/*-----*/
register admin support support contact='System Support'
grant authority support classes=system

/*-----*/
/* System Reporting Administrators */
/*-----*/
register admin reporter reporter contact='System Reporting'
```

```

/*-----*/
/* Client Administrators */
/*-----*/
register admin helpdesk helpdesk contact='Client Administrator'
/* grant authority helpdesk classes=node node=* */
/* You first have to register client nodes for this command to succeed! */

/*-----*/
/* Query all administrators */
/*-----*/
query admin

```

B.1.2 Define Client Option Sets

The following is an ADSM macro which contains the administrative commands to define our client option sets. The name of the macro is `mac.optionsets`.

```

/*=====*/
/* ADSM V3 Redbook Macro - Define Client Option sets */
/* Getting Started with ADSM - Implementation (SG24-5416) */
/* */
/* These scripts are supplied to help you implement an ADSM */
/* environment. The names and values reflect the recommendations */
/* for our Redbook environment. You may change them to suit your */
/* environment as required. */
/* */
/* This file is designed to be run as an Administrative Command */
/* Line macro. You should use the -itemcommit parameter when you */
/* start the admin session, otherwise the macro could fail!! */
/* */
/*=====*/

/*--Special Note-----*/
/* Remove the comments from the following delete cloptset */
/* commands if you want to rerun this macro. */
/*-----*/
/* delete cloptset redbook */
/* delete cloptset aix */
/* delete cloptset netware */
/* delete cloptset windows */

/* Base Option Set */
define cloptset redbook description="Redbook Base Set"
define clientopt redbook changingretries 4
define clientopt redbook compressalways yes
define clientopt redbook compression off
define clientopt redbook dirmc directory
define clientopt redbook domain all-local
define clientopt redbook maxcmdretries 2
define clientopt redbook memoryefficientbackup no
define clientopt redbook quiet
define clientopt redbook retryperiod 20
define clientopt redbook runasservice yes
define clientopt redbook schedmode prompted
define clientopt redbook scrolllines 20
define clientopt redbook scrollprompt no
define clientopt redbook subdir no
define clientopt redbook tapeprompt no
define clientopt redbook txmbytelimit 25600

/*-----*/
/* AIX Options Set */
/*-----*/
copy cloptset redbook aix
update cloptset aix description="AIX Clients"
define clientopt aix inclexcl "exclude /unix/" seq=1
define clientopt aix inclexcl "exclude.dir /unix/" seq=2
define clientopt aix inclexcl "exclude ../core" seq=3
define clientopt aix inclexcl "exclude /tmp/.../*" seq=4
define clientopt aix inclexcl "include ../dsmwebcl.log special" seq=5
define clientopt aix inclexcl "include ../dsmsched.log special" seq=6
define clientopt aix inclexcl "include ../dsmerror.log special" seq=7

```

```

/*-----*/
/* Netware Options Set */
/*-----*/
copy cloptset redbook netware
update cloptset netware description="Netware Clients"
define clientopt netware inclexcl "exclude sys:vol$log.err" seq=1
define clientopt netware inclexcl "exclude sys:tts$log.err" seq=2
define clientopt netware inclexcl "exclude sys:system/sys$log.err" seq=3
define clientopt netware inclexcl "exclude sys:system/events.log" seq=4
define clientopt netware inclexcl "exclude sys:system/secaudit.log" seq=5
define clientopt netware inclexcl "exclude sys:system/system.log" seq=6
define clientopt netware inclexcl "exclude sys:system/cmaster.dba" seq=7
define clientopt netware inclexcl "exclude sys:system/btrieve.trn" seq=8
define clientopt netware inclexcl "exclude sys:system/tsa/tsa$temp.*" seq=9
define clientopt netware inclexcl "include *:/.../dsmwebcl.log special" seq=10
define clientopt netware inclexcl "include *:/.../dmsched.log special" seq=11
define clientopt netware inclexcl "include *:/.../dsmerror.log special" seq=12

/*-----*/
/* Windows Options Set */
/*-----*/
copy cloptset redbook windows
update cloptset windows description="Windows Clients"
define clientopt windows inclexcl "exclude *:/.../pagefile.sys" seq=1
define clientopt windows inclexcl "exclude *:/.../netlogon.chg" seq=2
define clientopt windows inclexcl "exclude *:/.../system32/config/.../*" seq=3
define clientopt windows inclexcl "exclude *:/.../ntuser.dat" seq=4
define clientopt windows inclexcl "exclude *:/.../ntuser.dat.log" seq=5
define clientopt windows inclexcl "exclude *:/.../temp/.../*" seq=6
define clientopt windows inclexcl "exclude *:/.../cache/*" seq=7
define clientopt windows inclexcl "exclude *:/.../recycler/*" seq=8
define clientopt windows inclexcl "exclude *:/.../Temporary Internet Files/*" seq=9
define clientopt windows inclexcl "exclude *:/microsoft uam volume/.../*" seq=10
define clientopt windows inclexcl "exclude *:/ibmio.com" seq=11
define clientopt windows inclexcl "exclude *:/ibmdos.com" seq=12
define clientopt windows inclexcl "exclude *:/msdos.sys" seq=13
define clientopt windows inclexcl "exclude *:/io.sys" seq=14
define clientopt windows inclexcl "include *:/.../dsmwebcl.log special" seq=15
define clientopt windows inclexcl "include *:/.../dmsched.log special" seq=16
define clientopt windows inclexcl "include *:/.../dsmerror.log special" seq=17

/*-----*/
/* Query all client option sets */
/*-----*/
select * from cloptsets

```

B.1.3 Define Policy Structure

The following is an ADSM macro which contains the administrative commands to delete the default storage pools. The name of the macro is `mac.policy`.

```

/*=====*/
/* ADSM V3 Redbook Macro - Define Policy Domain,Set,Mgmt Class */
/* Getting Started with ADSM - Implementation (SG24-5416) */
/* */
/* These scripts are supplied to help you implement an ADSM */
/* environment. The names and values reflect the recommendations */
/* for our Redbook environment. You may change them to suit your */
/* environment as required. */
/* */
/* This file is designed to be run as an Administrative Command */
/* Line macro. You should use the -itemcommit parameter when you */
/* start the admin session, otherwise the macro could fail!! */
/* */
/*=====*/

/*-----*/
/* Policy Domains */
/*-----*/
define domain server description="Server nodes" backretention=100 \
archretention=365
define domain workstn description="Workstation nodes" backretention=100 \
archretention=365

```

```

/*-----*/
/* Policy Sets */
/*-----*/
define policyset server server description="Server nodes"
define policyset workstn workstn description="Workstation nodes"

/*-----*/
/* Management Classes */
/*-----*/
define mgmtclass server server data migdestination=NONE \
description="Default management class for server domain"
assign defmgmtclass server server data
define mgmtclass server server directory migdestination=NONE \
description="Directory management class for server domain"
define mgmtclass server server special migdestination=NONE \
description="Special management class for server domain"
define mgmtclass workstn workstn data migdestination=NONE \
description="Default management class for workstn domain"
assign defmgmtclass workstn workstn data
define mgmtclass workstn workstn directory migdestination=NONE \
description="Directory management class for workstn domain"
define mgmtclass workstn workstn special migdestination=NONE \
description="Special management class for workstn domain"

/*-----*/
/* Copy Groups */
/*-----*/
define copygroup server server data type=Backup destination=DISKDATA \
frequency=1 verexists=3 verdeleted=1 retextra=100 retonly=365 \
mode=modified serialization=shrstatic
define copygroup server server directory type=Backup destination=DISKDIRS \
frequency=1 verexists=nolimit verdeleted=1 retextra=100 retonly=365 \
mode=modified serialization=shrstatic
define copygroup server server special type=Backup destination=DISKDATA \
frequency=1 verexists=3 verdeleted=1 retextra=100 retonly=365 \
mode=modified serialization=shrdynamic
define copygroup workstn workstn data type=Backup destination=DISKDATA \
frequency=1 verexists=2 verdeleted=1 retextra=30 retonly=100 \
mode=modified serialization=shrstatic
define copygroup workstn workstn directory type=Backup destination=DISKDIRS \
frequency=1 verexists=nolimit verdeleted=1 retextra=30 retonly=100 \
mode=modified serialization=shrstatic
define copygroup workstn workstn special type=Backup destination=DISKDATA \
frequency=1 verexists=2 verdeleted=1 retextra=30 retonly=100 \
mode=modified serialization=shrdynamic
define copygroup server server data type=Archive destination=DISKDATA \
retver=365 serialization=shrstatic
define copygroup workstn workstn data type=Archive destination=DISKDATA \
retver=100 serialization=shrstatic

/*-----*/
/* Validate policy sets */
/*-----*/
validate policyset server server
validate policyset workstn workstn

```

B.1.4 Define Schedules

The following is an ADSM macro which contains the administrative commands to define our schedules. The name of the macro is `mac.schedules`.

```

/*=====*/
/* ADSM V3 Redbook Macro - Define Administrative/Client Schedules */
/* Getting Started with ADSM - Implementation (SG24-5416) */
/* */
/* These scripts are supplied to help you implement an ADSM */
/* environment. The names and values reflect the recommendations */
/* for our Redbook environment. You may change them to suit your */
/* environment as required. */
/* */
/* This file is designed to be run as an Administrative Command */
/* Line macro. You should use the -itemcommit parameter when you */
/* start the admin session, otherwise the macro could fail!! */
/* */
/*=====*/

```

```

/*-----*/
/* Offsite Backup Schedule */
/*-----*/
define schedule redbook_offsite type=admin cmd="run redbook_offsite" \
description="Backup all data for offsite storage" starttime=04:00 active=yes

/*-----*/
/* Volume History File Management */
/*-----*/
define schedule delete_volhist type=admin \
cmd="delete volhistory type=dbbackup todate=today-5" \
description="Delete volume history information for database backups" \
starttime=07:00 active=yes
define schedule backup_volhist type=admin cmd="backup volhistory" \
description="Backup volume history file" starttime=07:05 active=yes

/*-----*/
/* Disk Migration */
/*-----*/
define schedule MIGRATION_START type=admin cmd="update stgpool diskdata hi=0 lo=0"
description="Start migration on DISKDATA storage pool" starttime=07:00 active=yes
define schedule MIGRATION_STOP type=admin cmd="update stgpool diskdata hi=70 lo=30"
description="Stop migration on DISKDATA storage pool" starttime=10:00 active=yes

/*-----*/
/* Tape Space Reclamation */
/*-----*/
define schedule RECLAIM_OFFDIRS_START type=admin \
description="Start reclaim on the OFFDIRS storage pool" \
cmd="update stg offdirs rec=75" starttime=10:00 active=yes
define schedule RECLAIM_OFFDIRS_STOP type=admin \
description="Stop reclaim on the OFFDIRS storage pool" \
cmd="update stg offdirs rec=100" starttime=11:00 active=yes
define schedule RECLAIM_OFFDATA_START type=admin \
description="Start reclaim on the OFFDATA storage pool" \
cmd="update stg offdata rec=75" starttime=11:00 active=yes
define schedule RECLAIM_OFFDATA_STOP type=admin \
description="Stop reclaim on the OFFDATA storage pool" \
cmd="update stg offdata rec=100" starttime=14:00 active=yes
define schedule RECLAIM_TAPEDATA_START type=admin \
description="Start reclaim on the TAPEDATA storage pool" \
cmd="update stg tapedata rec=75" starttime=14:00 active=yes
define schedule RECLAIM_TAPEDATA_STOP type=admin \
description="Stop reclaim on the TAPEDATA storage pool" \
cmd="update stg tapedata rec=100" starttime=17:00 active=yes

/*-----*/
/* Inventory Expiration */
/*-----*/
define schedule EXPIRE_INVENTORY type=admin description="Inventory expiration" \
cmd="expire inventory" starttime=17:00 active=yes

/*-----*/
/* Audit Licences */
/*-----*/
define schedule AUDIT_LICENSE type=admin description="Audit licenses" \
cmd="audit licenses" starttime=00:00 active=yes

/*-----*/
/* Client Schedules */
/*-----*/
define schedule server server_nightly starttime=22:00 duration=3 durunits=hours \
description="Nightly backup schedule for SERVER domain"
define schedule workstn workstn_nightly starttime=22:00 duration=3 durunits=hours \
description="Nightly backup schedule for WORKSTN domain"

/*-----*/
/* Query all schedules */
/*-----*/
query schedule
query schedule type=admin

```

B.1.5 Define Server Scripts

The following is an ADSM macro which contains the administrative commands to define our server scripts. The name of the macro is `mac.scripts`.

```

/*=====*/
/* ADSM V3 Redbook Macro - Define Server Scripts */
/* Getting Started with ADSM - Implementation (SG24-5416) */
/* */
/* These scripts are supplied to help you implement an ADSM */
/* environment. The names and values reflect the recommendations */
/* for our Redbook environment. You may change them to suit your */
/* environment as required. */
/* */
/* This file is designed to be run as an Administrative Command */
/* Line macro. You should use the -itemcommit parameter when you */
/* start the admin session, otherwise the macro could fail!! */
/* */
/*=====*/

/*-----*/
/* Define Offsite Storage Process */
/*-----*/

delete script redbook_offsite
define script redbook_offsite description="Backup all data for offsite storage"
update script redbook_offsite /*-----*/
update script redbook_offsite /* Script Name: redbook_offsite */
update script redbook_offsite /* Description: Backup data for offsite storage */
update script redbook_offsite /*-----*/
update script redbook_offsite /* If there are active node sessions, then */
update script redbook_offsite /* reschedule this script to run again in 20 */
update script redbook_offsite /* 20 minutes. If there are no active node */
update script redbook_offsite /* sessions, then backup all the onsite storage */
update script redbook_offsite /* pools and the ADSM database. */
update script redbook_offsite /*-----*/

update script redbook_offsite "select * from sessions where -"
update script redbook_offsite "    upper(session_type)='NODE'"
update script redbook_offsite "if (rc_ok) goto reschedule"
update script redbook_offsite "backup stgpool diskdirs offdirs wait=yes"
update script redbook_offsite "backup stgpool diskdata offdata wait=yes"
update script redbook_offsite "backup stgpool tapedata offdata wait=yes"
update script redbook_offsite "backup db devclass=offsite type=full scratch=yes"
update script redbook_offsite "exit"
update script redbook_offsite "reschedule:"
update script redbook_offsite "delete schedule redbook_offsite_retry type=admin"
update script redbook_offsite "define schedule redbook_offsite_retry type=admin -"
update script redbook_offsite "    cmd='run redbook_offsite' active=yes -"
update script redbook_offsite "    startt=NOW+0:20 perunits=onetime"

/*-----*/
/* Query all redbook scripts */
/*-----*/
query script redbook*

```

B.1.6 Create Storage Pools

The following is an ADSM macro which contains the administrative commands to delete the default storage pools. The name of the macro is `mac.stgcreate`.

```

/*=====*/
/* ADSM V3 Redbook Macro - Create Storage Pools */
/* Getting Started with ADSM - Implementation (SG24-5416) */
/* */
/* These scripts are supplied to help you implement an ADSM */
/* environment. Names for device classes used should be entered */
/* between the << >> symbols and these << >> symbols removed. */
/* The MAXScratch is set purposefully at a high value to avoid */
/* any misleading out of space messages. You may set it to a more */
/* meaningful value to better reflect the %util value of your */
/* pool when q stgpool is run. */
/* */
/* This file is designed to be run as an Administrative Command */
/* Line macro. You should use the -itemcommit parameter when you */
/* start the admin session, otherwise the macro could fail!! */
/* */
/*=====*/

```

```

define stgpool tapedata <<library device class name>> highmig=100 \
maxscratch=10000 collocate=no reclaim=100 reusedelay=1
define stgpool diskdata disk nextstgpool=tapedata highmig=70 lowmig=30 \
cache=no
define stgpool diskdirs disk highmig=100
define stgpool offdirs <<offsite library device class name>> pooltype=copy \
reclaim=100 maxscratch=10000 reusedelay=5
define stgpool offdata <<offsite library device class name>> pooltype=copy \
reclaim=100 maxscratch=10000 reusedelay=5
define stgpool none disk

```

B.1.7 Delete Default Storage Pools

The following is an ADSM macro which contains the administrative commands to delete the default storage pools. The name of the macro is `mac.stgdelete`.

```

/*=====*/
/* ADSM V3 Redbook Macro - Delete Default Storage Pools */
/* Getting Started with ADSM - Implementation (SG24-5416) */
/* */
/* These scripts are supplied to help you implement an ADSM */
/* environment. The names and values reflect the recommendations */
/* for our Redbook environment. You may change them to suit your */
/* environment as required. */
/* */
/* This file is designed to be run as an Administrative Command */
/* Line macro. You should use the -itemcommit parameter when you */
/* start the admin session, otherwise the macro could fail!! */
/* */
/*=====*/

delete stgpool backuappool
delete stgpool archivepool
delete stgpool spacemgpool
delete stgpool diskpool

```

B.2 Server Options Files

We have created server option files for the AIX, MVS, and Windows NT server environments. Our environment assumes that TCP/IP is the network protocol, Web access is enabled for administrators, and basic performance tuning values are specified.

All possible server options for each platform have been specified in the respective files. Server options have been grouped into the same categories as they appear in the administrator reference manuals for ease of reference. Those categories are:

- Communications
- Automated Cartridge System Library
- Client-server communication processing
- Site dependent
- Database and recovery log
- Group and transfer data
- Messages
- Event logging
- Miscellaneous

Within each category, the options are ordered alphabetically. Options which are not active are preceded by an "*". All other options are active. Inactive options are those that are either not applicable for our TCPIP assumption or require further setup such as those for Tivoli event monitoring.

B.2.1 AIX

The following shows our server options file for AIX. This file can also be used as the basis for a server options file for the HP/UX and Solaris platforms.

ADSM HP/UX and Solaris servers only support TCPIP as a network protocol. HP/UX does not support the STK Automated Cartridge System Library Software or the ENABLE3590LIBRARY parameters.

```

=====
* ADSM V3 Redbook Server Options File - AIX Version *
* Getting Started with ADSM - Implementation (SG24-5416) *
=====

*-----*
*Communications *
*-----*
COMMETHOD HTTP
COMMETHOD SHAREDMEM
COMMETHOD TCPIP
HTTPPORT 1580
HTTPSPORT 1543
*IPXBUFFERSIZE
*IPXSOCKET
*LANADAPTER
*LUNAME
*NETBIOSBUFFERSIZE
*NETBIOSNAME
*NETBIOSSESSION
SHMPORT 1510
*SNMPHEARTBEATINTERVAL
*SNMPMESSAGECATEGORY
*SNMPSUBAGENT
*SNMPSUBAGENTHOST
*SNMPSUBAGENTPORT
TCPBUFSIZE 32
TCPNODELAY YES
TCPPORT 1500
TCPWINDOWSIZE 2048
*TFNAME
*TFNPROFILENAME

*-----*
*Automated Cartridge System Library Software *
*-----*
*ACACCESSID
*ACSLCKDRIVE
*ACSQUICKINIT
*ACSTIMEOUTX

*-----*
* Client-Server *
*-----*
COMMTIMEOUT 60
DISABLESCHEDS YES
IDLETIMEOUT 15
MAXSESSIONS 25
USELARGEBUFFERS YES

*-----*
* Site Dependent *
*-----*
DATEFORMAT 2
LANGUAGE AMENG
NUMBERFORMAT 1
TIMEFORMAT 1

```

```

*-----*
* Database & Recovery Log *
*-----*
BUFFPOOLSIZE 32768
LOGPOOLSIZE 2048
MIRRORREAD DB NORMAL
MIRRORREAD LOG NORMAL
MIRRORWRITE DB SEQUENTIAL
MIRRORWRITE LOG SEQUENTIAL

*-----*
* Group *
*-----*
MOVEBATCHSIZE 500
MOVESIZETHRESH 256
TXNGROUPMAX 256

*-----*
* Message Options *
*-----*
EXPQUIET YES
MESSAGEFORMAT 1
MSGINTERVAL 1
STATUSMSGCNT 10

*-----*
* Event Logging *
*-----*
EVENTSERVER YES
FILEEXIT NO
USEREXIT NO
*TECBEGINEVENTLOGGING
*TECHOSTNAME
*TECPORT

*-----*
* Miscellaneous *
*-----*
DEVCONFIG /adsm/devconfig
DEVCONFIG /adsm_mirror/devconfig
ENABLE3590LIBRARY NO
EXPINTERVAL 0
*NOAUDITSTORAGE
*NOPREEMPT
RESTORE INTERVAL 1440
VOLUMEHISTORY /adsm/volumehistory
VOLUMEHISTORY /adsm_mirror/volumehistory

```

B.2.2 MVS

The following shows our server options file for MVS.

```

*=====*
* ADSM V3 Redbook Server Options File - MVS Version *
* Getting Started with ADSM - Implementation (SG24-5416) *
*=====*

*-----*
* Communications *
*-----*
*CLIOPORT
*HTIPICSPORT
HTIPTCPORT 1580
*ICSSNAME
*ICSPORT
*IUCV
*LUNAME
TCPBUFSIZE 32
TCPNAME TCPIP
TCPPORT 1500
*TPNPROFILENAME

```

```

*-----*
* Client-Server *
*-----*
COMMTIMEOUT 60
DISABLESCHEDS YES
IDLETIMEOUT 15
MAXSESSIONS 25
USELARGEBUFFERS YES

*-----*
* Site Dependent *
*-----*
DATEFORMAT 2
LANGUAGE AMENG
NUMBERFORMAT 1
TIMEFORMAT 1

*-----*
* Database & Recovery Log *
*-----*
BUFPoolsize 16384
LOGPoolsize 2048
MIRRORRead DB Normal
MIRRORRead LOG Normal
MIRRORWrite DB Sequential
MIRRORWrite LOG Sequential

*-----*
* Group *
*-----*
MOVEBatchsize 500
MOVESizethresh 256
TXNGroupmax 256

*-----*
* Message Options *
*-----*
EXPQIet YES
MESsageformat 1
MSGHIGHLIGHT 0
MSGSUPPRESS 0
ROUTECode 11

*-----*
* Event Logging *
*-----*
EVENTSERVer YES
FILEEXIT NO
USEREXIT NO
*TECBeginEventLogging
*TECHostname
*TECPort

*-----*
* Miscellaneous *
*-----*
DELetionexit ARCTVEXT
DEVCONFIG 'ADSM.SERVER.DEVCONF'
DEVCONFIG 'ADSM.SERVER.DEVCONF.ALT'
EXPInterval 0
*HPDTENABLE
*HPDTSTORAGE
*HPDTBUFFER
*HPDTEXPAND
LICENSE CLIENTS 1
*NOAUDITSTORAGE
*NOPREEMPT
RESTORE INTERVAL 1440
VOLUMEHistory 'ADSM.SERVER.VOLHIST'
VOLUMEHistory 'ADSM.SERVER.VOLHIST.ALT'

```

B.2.3 Windows NT

The following shows our server options file for Windows NT.

```

=====
* ADSM V3 Redbook Server Options File - Windows NT Version *
* Getting Started with ADSM - Implementation (SG24-5416) *
=====

*-----*
*Communications *
*-----*
COMMETHOD HTTP
COMMETHOD NAMEDPIPE
COMMETHOD TCP/IP
HTTPPORT 1580
HTTPPORT 1543
*IPXBUFFERSIZE
*IPXSOCKET
*LANADAPTER
NAMEDPIPE\NAME \\.\PIPE\ADSMPIPE
*NETBIOSBUFFERSIZE
*NETBIOSNAME
*NETBIOSSESSION
*SNMPHEARTBEATINTERVAL
*SNMPMESSAGECATEGORY
*SNMPSUBAGENT
*SNMPSUBAGENTHOST
*SNMPSUBAGENTPORT
TCPBUFSIZE 32
TCPNODELAY YES
TCPPORT 1500
TCPWINDOWSIZE 2048

*-----*
* Client-Server *
*-----*
COMTIMEOUT 60
DISABLESCHEDULES YES
IDLETIMEOUT 15
MAXSESSIONS 25
USELARGEBUFFERS YES

*-----*
* Site Dependent *
*-----*
DATEFORMAT 2
LANGUAGE AMENG
NUMBERFORMAT 1
TIMEFORMAT 1

*-----*
* Database & Recovery Log *
*-----*
BUFFPOOLSIZE 16384
LOGPOOLSIZE 2048
MIRRORREAD DB NORMAL
MIRRORREAD LOG NORMAL
MIRRORWRITE DB SEQUENTIAL
MIRRORWRITE LOG SEQUENTIAL

*-----*
* Group *
*-----*
MOVEBATCHSIZE 500
MOVESIZETHRESH 256
TXNGROUPMAX 256

*-----*
* Message Options *
*-----*
EXPQULET YES
MESSAGEFORMAT 1
MSGINTERVAL 1
STATUSMSGCNT 10

```

```

*-----*
* Event Logging *
*-----*
EVENTSERVER YES
FILEEXIT NO
USEREXIT NO
*TECBEGINEVENTLOGGING
*TECHOSTNAME
*TECPORT

*-----*
* Miscellaneous *
*-----*
DEVCONFIG c:\adsm\devconfig
DEVCONFIG c:\adsm_mirror\devconfig
EXPINTERVAL 0
*NOAUDITSTORAGE
*NOPREEMPT
RESTORE INTERVAL 1440
VOLUMEHISTORY c:\adsm\volumehistory
VOLUMEHISTORY c:\adsm_mirror\volumehistory

```

B.3 Client Options Files

We have created client option files for the NetWare and Windows 95/98/NT environments. Our environment assumes that TCPIP is the network protocol, web access is enabled for administrators, and basic performance tuning values are specified. You must replace fields surrounded by angle brackets (<< and >>) with the specific values for your site.

B.3.1 AIX

B.3.1.1 Client System Options File

The following shows our client system options file for AIX. This file can also be used as the basis for a client system options file for the HP/UX and Solaris platforms.

```

*-----*
* AD SM V3 Redbook Client System Options File (dsm.sys)- AIX Version *
* Getting Started with ADSM - Implementation (SG24-5416) *
*-----*

*-----*
* Communications *
*-----*
SERVERNAME <<ADSM-Server-Name>>
COMMMETHOD TCPIP
TCPBUFFSIZE 32
TCPNODELAY Yes
TCPPORT 1500
TCPSEVERADDRESS <<ADSM-Server-Address>>
TCPWINDOWSIZE 58
LARGECOMMBUFFERS Yes

*-----*
* Operations *
*-----*
*ERRORLOGNAME
ERRORLOGRETENTION 10,D
PASSWORDACCESS GENERATE
SCHEDLOGNAME dsm sched.log
SCHEDLOGRETENTION 10,D

```

```

*-----*
* Site Dependent *
*-----*
***These site dependent options are supported on all UNIX
*DATEFORMAT 1
*LANGUAGE AMENG
*NUMBERFORMAT 1
*TIMEFORMAT 1

*-----*
* Include/Exclude *
*-----*
*None - specified via the Client Option Set (AIX) instead

```

B.3.1.2 Client User Options File

The following shows our client user options file for AIX. This file can also be used as the basis for a client user options file for the HP/UX and Solaris platforms.

```

*-----*
* ADSM V3 Redbook Client User Options File (dsm.opt)- AIX Version *
* Getting Started with ADSM - Implementation (SG24-5416) *
*-----*

*-----*
* Communications *
*-----*
SERVERNAME <<ADSM-Server-Name>>

```

B.3.2 NetWare

The following shows our client options file for NetWare.

```

*-----*
* ADSM V3 Redbook Client Options File - Netware Version *
* Getting Started with ADSM - Implementation (SG24-5416) *
*-----*

*-----*
* Identification *
*-----*
NODENAME <<Netware-Server-Name>>

*-----*
* Communications *
*-----*
COMMETHOD TCP/IP
TCPBUFFSIZE 63
TCPNODELAY Yes
TCPPOST 1500
TCPSEVERADDRESS <<ADSM-Server-Name>>
TCPWINDOWSIZE 32

*-----*
* Operations *
*-----*
*ERRORLOGNAME
ERRORLOGRETENTION 10,D
PASSWORDACCESS GENERATE
REPLACE PROMPT
SCHEDLOGNAME dmsched.log
SCHEDLOGRETENTION 10,D

```

```

*-----*
* Site Dependent *
*-----*
DATEFORMAT 1
LANGUAGE AMENG
NUMBERFORMAT 1
TIMEFORMAT 1

*-----*
* Include/Exclude *
*-----*
*None - specified via the Client Option Set (NETWARE) instead

```

B.3.3 Windows 95/98/NT

The following shows our client options file for Windows 95/98/NT.

```

=====
* ADSM V3 Redbook Client Options File - Windows Version *
* Getting Started with ADSM - Implementation (SG24-5416) *
=====

*-----*
* Communications *
*-----*
COMMMETHOD TCPIP
TCPBUFFSIZE 63
TCPNODELAY Yes
TCPPORT 1500
TCPSEVERADDRESS <<ADSM-Server-Name>>
TCPWINDOWSIZE 32

*-----*
* Operations *
*-----*
*ERRORLOGNAME
ERRORLOGRETENTION 10,D
PASSWORDACCESS GENERATE
REPLACE PROMPT
SCHEDLOGNAME dmsched.log
SCHEDLOGRETENTION 10,D

*-----*
* Site Dependent *
*-----*
DATEFORMAT 1
LANGUAGE AMENG
NUMBERFORMAT 1
TIMEFORMAT 1

*-----*
* Include/Exclude *
*-----*
*None - specified via the Client Option Set (WINDOWS) instead

```

Appendix C. Special Notices

This publication is intended to help IBM customers, Business Partners, consultants, and staff implement ADSTAR Distributed Storage Manager (ADSM). The information in this publication is not intended as the specification of any programming interfaces that are provided by ADSM. See the PUBLICATIONS section of the IBM Programming Announcement for ADSM for more information about what publications are considered to be product documentation.

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 IBM's product, program, or service may be used. Any functionally equivalent program that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program or service.

Information in this book was developed in conjunction with use of the equipment specified, and is limited in application to those specific hardware and software products and levels.

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, NY 10594 USA.

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, Dept. 600A, Mail Drop 1329, Somers, NY 10589 USA.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

Any pointers in this publication to external Web sites are provided for convenience only and do not in any manner serve as an endorsement of these Web sites.

Any performance data contained in this document was determined in a controlled environment, and therefore, the results that may be obtained in other operating environments may vary significantly. Users of this document should verify the applicable data for their specific environment.

Reference to PTF numbers that have not been released through the normal distribution process does not imply general availability. The purpose of including

these reference numbers is to alert IBM customers to specific information relative to the implementation of the PTF when it becomes available to each customer according to the normal IBM PTF distribution process.

The following terms are trademarks of the International Business Machines Corporation in the United States and/or other countries:

IBM ®	ADSTAR
AIX	AS/400
MVS/ESA	NetView
OS/2	OS/390
OS/400	

The following terms are trademarks of other companies:

C-bus is a trademark of Corollary, Inc. in the United States and/or other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and/or other countries.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States and/or other countries.

PC Direct is a trademark of Ziff Communications Company in the United States and/or other countries and is used by IBM Corporation under license.

ActionMedia, LANDesk, MMX, Pentium and ProShare are trademarks of Intel Corporation in the United States and/or other countries. (For a complete list of Intel trademarks see www.intel.com/tradmarx.htm.)

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

Other company, product, and service names may be trademarks or service marks of others.

Appendix D. Related Publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

D.1 International Technical Support Organization Publications

For information on ordering these ITSO publications see "How to Get ITSO Redbooks" on page 263.

ADSM Redbooks

Book Title	Publication Number
General Topics	
ADSM Version 3 Technical Guide	SG24-2236
ADSM Concepts	SG24-4877
ADSM Advanced Implementation Experiences	GG24-4221
Using ADSM Hierarchical Storage Management	SG24-4631
Client Disaster Recovery: Bare Metal Restore	SG24-4880
Specific Server Books	
Windows NT Backup and Recovery with ADSM	SG24-2231
ADSM Server-to-Server Implementation and Operation	SG24-5244
ADSM Server for Windows NT Configuration and Recovery Examples	SG24-4878
Getting Started with ADSM/6000	GG24-4421
ADSM for AIX: Advanced Topics	SG24-4601
AIX Tape Management	SG24-4705
ADSM/6000 on 9076 SP2	GG24-4499
ADSM for MVS: Recovery and Disaster Recovery	SG24-4537
ADSM for MVS: Using Tapes and Tape Libraries	SG24-4538
Getting Started with ADSM/2	GG24-4321
ADSM for OS/2: Advanced Topics	SG24-4740
Setting Up and Implementing ADSM/400	GG24-4460
ADSM/VSE Implementation Guide	SG24-4266
Specific Client Books	
Getting Started with ADSM NetWare Clients	GG24-4242
Getting Started with ADSM AIX Clients	GG24-4243
ADSM API Examples for OS/2 and Windows	SG24-2588
ADSM with Other Products	
A Practical Guide to Network Storage Manager	SG24-2242

Book Title	Publication Number
Using ADSM to Back Up Databases	SG24-4335
Using ADSM to Back Up Lotus Notes	SG24-4534
Using ADSM to Back Up and Recover Microsoft Exchange Server	SG24-5266
Hierarchical Storage Management for NetWare: ADSM and AvailHSM Implementation	SG24-4713
Using ADSM to Back Up OS/2 LAN Server and Warp Server	SG24-4682
Backup, Recovery, and Availability with DB2 Parallel Edition on RISC/6000	SG24-4695
ADSM Operation and Management with TME10	SG24-2214
ADSM Reporting with SAMS:Vantage	SG24-5271

D.2 Redbooks on CD-ROMs

Redbooks are also available on CD-ROMs. **Order a subscription** and receive updates 2-4 times a year.

CD-ROM Title	Subscription Number	Collection Kit Number
System/390 Redbooks Collection	SBOF-7201	SK2T-2177
Networking and Systems Management Redbooks Collection	SBOF-7370	SK2T-6022
Transaction Processing and Data Management Redbook	SBOF-7240	SK2T-8038
Lotus Redbooks Collection	SBOF-6899	SK2T-8039
Tivoli Redbooks Collection	SBOF-6898	SK2T-8044
AS/400 Redbooks Collection	SBOF-7270	SK2T-2849
RS/6000 Redbooks Collection (HTML, BkMgr)	SBOF-7230	SK2T-8040
RS/6000 Redbooks Collection (PostScript)	SBOF-7205	SK2T-8041
RS/6000 Redbooks Collection (PDF Format)	SBOF-8700	SK2T-8043
Application Development Redbooks Collection	SBOF-7290	SK2T-8037

D.3 ADSM Product Publications

Book Title	Publication Number
ADSM V3R1 Messages	GC35-0271
ADSM V3R1 AIX Quick Start	GC35-0273
ADSM V3R1 AIX Administrator's Guide	GC35-0274
ADSM V3R1 AIX Administrator's Reference	GC35-0275
ADSM V3R1 MVS Quick Start	GC35-0276

Book Title	Publication Number
ADSM V3R1 MVS Administrator's Guide	GC35-0277
ADSM V3R1 MVS Administrator's Reference	GC35-0278

Book Title	Publication Number
ADSM V3R1 AIX License	SC35-0283
ADSM V3R1 MVS License	GC35-0284
ADSM V3R1 Windows NT Administrator's Guide	GC35-0292
ADSM V3R1 Windows NT Administrator's Reference	GC35-0293
ADSM V3R1 Windows NT License	SC35-0294
ADSM V3R1 Windows NT Quick Start	GC35-0295
ADSM V3R1 Using the UNIX Backup-Archive Clients	SH26-4075
ADSM V3R1 Using the OS/2 Backup-Archive Client	SH26-4076
ADSM V3R1 Using the Novell NetWare Backup-Archive Client	SH26-4077
ADSM V3R1 Using the Microsoft Windows Backup-Archive Client	SH26-4078
ADSM V3R1 Installing the Clients	SH26-4080
ADSM V3R1 Using the Application Programming Interface	SH26-4081
ADSM V3R1 Trace Facility Guide	SH26-4082
ADSM V3R1 Client Reference Cards	SX26-6019
ADSM V3R1 Using the OS/2 Lotus Notes Backup Agent	SH26-4084
ADSMConnect Agent for Oracle7 Backup on AIX Installation and User's Guide	SH26-4061
ADSMConnect Agent for Oracle Backup on Sun Solaris Installation and User's Guide	SH26-4063
ADSMConnect Agent for Lotus Notes on Windows NT Installation and User's Guide	SH26-4065
ADSMConnect Agent for Lotus Notes on AIX Installation and User's Guide	SH26-4067
ADSMConnect Agent for Microsoft SQL Server Installation and User's Guide	SH26-4069

D.4 ADSM Online Product Library

All of the ADSM publications are available in online readable format on the CD-ROM listed below.

CD-ROM Title	Publication Number
ADSM V3R1.0 MVS Online Product Library	SK3T-1396

How to Get ITSO Redbooks

This section explains how both customers and IBM employees can find out about ITSO redbooks, redpieces, and CD-ROMs. A form for ordering books and CD-ROMs by fax or e-mail is also provided.

- **Redbooks Web Site** <http://www.redbooks.ibm.com/>

Search for, view, download or order hardcopy/CD-ROM redbooks from the redbooks web site. Also read redpieces and download additional materials (code samples or diskette/CD-ROM images) from this redbooks site.

Redpieces are redbooks in progress; not all redbooks become redpieces and sometimes just a few chapters will be published this way. The intent is to get the information out much quicker than the formal publishing process allows.

- **E-mail Orders**

Send orders via e-mail including information from the redbooks fax order form to:

	e-mail address
In United States	usib6fpl@ibmmail.com
Outside North America	Contact information is in the "How to Order" section at this site: http://www.elink.ibm.ibm.com/pbl/pbl/

- **Telephone Orders**

United States (toll free)	1-800-879-2755
Canada (toll free)	1-800-IBM-4YOU
Outside North America	Country coordinator phone number is in the "How to Order" section at this site: http://www.elink.ibm.ibm.com/pbl/pbl/

- **Fax Orders**

United States (toll free)	1-800-445-9269
Canada	1-403-267-4455
Outside North America	Fax phone number is in the "How to Order" section at this site: http://www.elink.ibm.ibm.com/pbl/pbl/

This information was current at the time of publication, but is continually subject to change. The latest information for customer may be found at <http://www.redbooks.ibm.com/> and for IBM employees at <http://w3.itso.ibm.com/>.

IBM Intranet for Employees

IBM employees may register for information on workshops, residencies, and redbooks by accessing the IBM Intranet Web site at <http://w3.itso.ibm.com/> and clicking the ITSO Mailing List button. Look in the Materials repository for workshops, presentations, papers, and Web pages developed and written by the ITSO technical professionals; click the Additional Materials button. Employees may also view redbook, residency, and workshop announcements at <http://inews.ibm.com/>.

Index

A

access control lists

See ACLs

ACLs

backup of directory information

ACTIVATE POLICYSET command 121

ADMIN, admin ID 124

administrative client

ACTIVATE POLICYSET command 121

ADMIN ID 124

ASSIGN DEFMGMTCLASS command 109

AUDIT LICENSES command 140, 191

AUDIT VOLUME command 94

BACKUP DB command 62, 224

BACKUP STGPOOL command 95, 217, 220

BACKUP VOLHISTORY command 189

batch mode option of command line interface 146

BEGIN EVENTLOGGING command 203

CANCEL PROCESS command 201

CANCEL RESTORE command 201, 212

CANCEL SESSION command 178, 201

CHECKIN LIBVOLUME command 90, 96, 222, 225

CHECKOUT LIBVOLUME command 96, 219, 224

client option set

See client option set

command line interface 145

CONSOLE parameter 199

console session 199

CONTACT option 123

COPY CLOPTSET command 134

DEFINE ASSOCIATION command 192, 193

DEFINE CLIENTOPT command 134

DEFINE CLOPTSET command 134

DEFINE COPYGROUP command 112, 113

DEFINE DBBACKUPTRIGGER command 53

DEFINE DBCOPY command 58

DEFINE DBVOLUME command 43, 48

DEFINE DEVCLASS command 78, 87

DEFINE DOMAIN command 107

DEFINE DRIVE command 78, 86

DEFINE LIBRARY command 77, 85

DEFINE LOGCOPY command 59

DEFINE LOGVOLUME command 43, 50

DEFINE MGMTCLASS command 109

DEFINE POLICYSET command 108

DEFINE SCHEDULE command 185, 192

DEFINE SPACETRIGGER command 54

DEFINE STGPOOL command 81, 89

DEFINE VOLUME command 43, 92

DELETE CLIENTOPT command 134

DELETE CLOPTSET command 135

DELETE DBVOLUME command 60

DELETE DEVCLASS command 99

DELETE DOMAIN command 118

DELETE FILESPACE command 118, 130

DELETE LIBRARY command 99

DELETE LOGVOLUME command 61

DELETE STGPOOL command 93, 98

DELETE VOLHISTORY command 64, 188, 225

deleting client node filespaces 130

DISABLE SESSIONS command 200

DISMOUNT VOLUME command 94, 99

dsmadm program 145

ENABLE EVENTS command 203

EXPIRE INVENTORY command 121, 191, 232

EXTEND DB command 60, 63

EXTEND LOG command 61, 63

GRANT AUTHORITY command 123

granting node authority to administrator 129

GUI on Windows platform 143

HALT command 200

HELPDESK ID 125

installation 144

ITEMCOMMIT parameter 42

LABEL LIBVOLUME command 45, 83, 97, 215

loop mode option of command line interface 146

MACRO command 42

managing and controlling resources 143

MOVE DATA command 101

MOVE MEDIA command 208

privileges to manage resources 123

QUERY ACTLOG command 210, 226

QUERY ADMIN command 123, 126

QUERY AUDITOCUPANCY command 191, 207

QUERY CONTENT command 101

QUERY COPYGROUP command 119

QUERY DB command 204, 231

QUERY DBBACKUPTRIGGER command 53

QUERY DBVOLUME command 57, 205

QUERY DEVCLASS command 62, 87

QUERY DRIVE command 207

QUERY EVENT command 194, 211

QUERY FILESPACE command 130

QUERY LIBRARY command 62

QUERY LIBVOLUME command 92, 208

QUERY LICENSE command 139

QUERY LOG command 205, 233

QUERY LOGVOLUME command 58, 206

QUERY MEDIA command 208

QUERY NODE command 129

QUERY OCCUPANCY command 102, 208

QUERY OPTIONS command 38

QUERY PROCESS command 95, 131, 201, 212

QUERY REQUEST command 91, 97, 209

QUERY RESTORE command 201, 212

QUERY SCHEDULE command 193

QUERY SESSION command 213

QUERY STATUS command 40, 52, 200, 202

QUERY STGPOOL command 82, 98, 209

QUERY VOLHISTORY command 206, 224, 225

QUERY VOLUME command 190, 209, 218, 221

QUIESCE option 200

REDUCE DB command 64

REDUCE LOG command 64

REGISTER ADMIN command 123

- REGISTER LICENSE command 137
- REGISTER NODE command 127, 128
- registering admin IDs 124
- registering client nodes 128
- REMOVE ADMIN command 127
- REMOVE NODE command 118
- removing client nodes 131
- RENAME STGPOOL command 102
- REPLY command 91, 97
- RESET BUFPOOL command 232
- RESET LOGCONSUMPTION command 52
- scheduling backup of volume history 189
- scheduling deletion of volume history 188
- scheduling inventory expiration 191
- scheduling license audit 191
- scheduling offsite backup 186
- scheduling operations 184
- scheduling reclamation 190
- scheduling storage pool migration 189
- SELECT command 42, 103, 133, 188
- SERVER_CONSOLE ID 123, 196
- SET ACCOUNTING command 42
- SET ACTLOGRETENTION command 203
- SET AUTHENTICATION command 41
- SET INPWLENGTH command 41
- SET INVALIDPWLIMIT command 41
- SET LICENSEAUDITPERIOD command 140
- SET LOGMODE command 51
- SET PASSEXP command 41
- SET WEBAUTHTIMEOUT command 41
- TCPSERVERNAME option 145
- TSO interface on MVS 143
- UPDATE ADMIN command 126
- UPDATE NODE command 119, 130, 133
- UPDATE SPACETRIGGER command 54, 55
- UPDATE STGPOOL command 82, 98, 189, 190
- UPDATE VOLUME command 220, 222
- VALIDATE POLICYSET command 120
- verifying schedule events 211
- Web administrative interface
 - See Web administrative interface
- administrator
 - See administrative client
- ADSM Web page
 - list of supported storage devices 67
 - software fixes 37
- ARCHIVE command 163
- archive copy group
 - ARCHMC option 163
 - assigning files to archive copy group 163
 - defining a copy group 113
 - DESTINATION option 113
 - RETVR option 113
 - SERIALIZATION option 113
 - TYPE option 113
- archive description 163
- ARCHMC
 - ARCHIVE command 163
 - client option 108
- ARCHRET option, DEFINE DOMAIN command 107

- ASSIGN DEFMGMTCLASS command 109
- assigning files to specific management class 155, 163
- AUDIT LICENSES command 140, 191
- AUDIT VOLUME command 94

B

- BACKRET option, DEFINE DOMAIN command 107
- backup copy group
 - active file 110
 - assigning files to backup copy group 155
 - defining a copy group 110
 - DESTINATION option 111
 - expired file 110
 - FREQUENCY option 111
 - inactive file 110
 - INCLUDE option 155
 - MODE option 111
 - retention period 110
 - RETEXTRA option 110
 - REONLY option 111
 - SERIALIZATION option 112
 - TYPE option 111
 - VERDELETED option 110
 - VEREXISTS option 110
- BACKUP DB command 62, 224
- BACKUP STGPOOL command 95, 217, 220
- BACKUP VOLHISTORY command 189
- backup-archive client
 - ARCHIVE command 163
 - archive description 163
 - archiving files to server 163
 - ARCHMC option 108
 - assigning files to archive copy group 163
 - assigning files to backup copy group 155
 - associating client with schedule 193
 - automatic start of scheduler 177
 - backing up files to server 159
 - batch mode option of command line interface 158
 - central scheduling 174
 - changing node password 130
 - CHANGINGRETRIES option 112
 - CLIENT ID 128
 - client option set
 - See client option set
 - client system options file 153
 - client user options file 153
 - command line interface 157
 - COMMETHOD option 154
 - communication options 154
 - COMPRESSALWAYS option 235
 - COMPRESSION option 235
 - deleting files during archive operation 163
 - DIRMC option 108, 236
 - DOMAIN option 127
 - dsm program 167
 - dsm.opt file 153
 - dsm.sys file 153
 - DSM_CONFIG variable 152
 - DSM_DIR variable 152
 - DSM_LOG variable 152

- dsmc program 157
- DSMCUTIL program 177
- DSMTCA, client communication agent 152
- EXCLUDE option 155
- granting node authority to administrator 129
- graphical user interface 167
- INCLEXCL option 155
- INCLUDE option 108, 155
- include-exclude list 108, 155
- INCREMENTAL command 159
- LARGECOMMBUFFERS option 236
- loop mode option of command line interface 158
- manual start of scheduler 175
- NODENAME option 154
- operational options 155
- options file 153
- PASSWORDACCESS option 155, 156
- point-in-time restore 171
- Preferences file for Macintosh 153
- QUIET option 236
- QUIT command 166
- registering client nodes 128
- removing client nodes 131
- REPLACE option 155
- restartable restore sessions 171
- RESTORE command 160
- restoring backed up files to client machine 160
- restoring inactive file copies 161
- restoring most recent file copies 161
- RETRIEVE command 164
- retrieving archived files to client machine 164
- REVOKEREMOTEACCESS option
 - See Web client
- SCHEDMODE option 155
- SCHEDULE command 175
- scheduling client operation 192
- SELECTIVE command 159
- site dependent options 155
- SUBDIR option 155
- TAPEPROMPT option 155
- TCPBUFFSIZE option 154, 236
- TCPPORT option 154
- TCPSERVERADDRESS option 154
- TCPWINDOWSIZE option 154, 236
- TXNBYTELIMIT option 235, 237
- TXNGROUPMAX option 237
- USERID option 127
 - verifying schedule events 194
- banner frame, Web administrative interface 148
- batch mode of command line interface 146, 158
- BEGIN EVENTLOGGING command 203
- BUFPOOLSIZE, server option 40, 231

C

- CACHE option, DEFINE STGPOOL command 73
- CANCEL PROCESS command 201
- CANCEL RESTORE command 201, 212
- CANCEL SESSION command 178, 201
- CHANGINGRETRIES, client option 112
- charging of AD SM services 35

- CHECKIN LIBVOLUME command 90, 96, 222, 225
- CHECKLABEL option, DISMOUNT VOLUME command 94
- checklists
 - client implementation, client tasks 7
 - client implementation, server tasks 7
 - daily operations 8
 - planning 5
 - server implementation 5
- CHECKOUT LIBVOLUME command 96, 219, 224
- client acceptor, Web client 152, 180
- client access authority
 - See Web client
- client node
 - See backup-archive client
- client option set
 - adding options to set 134
 - associating client node 133
 - defining set 134
 - deleting set 135
 - easing administration of client nodes by centralized management 132
 - include-exclude statements 132
 - protecting option settings against overwriting 132
 - removing option from set 134
- client owner authority
 - See Web client
- CLIENT, client node ID 128
- COLLOCATE option, DEFINE STGPOOL command 73
- collocation of sequential access storage volumes 73
- command line client, Web administrative interface 147, 148
- commands, administrator
 - ACTIVATE POLICYSET 121
 - ASSIGN DEFMGMTCLASS 109
 - AUDIT LICENSES 140, 191
 - AUDIT VOLUME 94
 - BACKUP DB 62, 224
 - BACKUP STGPOOL 95, 217, 220
 - BACKUP VOLHISTORY 189
 - BEGIN EVENTLOGGING 203
 - CANCEL PROCESS 201
 - CANCEL RESTORE 201, 212
 - CANCEL SESSION 178, 201
 - CHECKIN LIBVOLUME 90, 96, 222, 225
 - CHECKOUT LIBVOLUME 96, 219, 224
 - COPY CLOPTSET 134
 - DEFINE ASSOCIATION 192, 193
 - DEFINE CLIENTOPT 134
 - DEFINE CLOPTSET 134
 - DEFINE COPYGROUP 112, 113
 - DEFINE DBBACKUPTRIGGER 53
 - DEFINE DBCOPY 58
 - DEFINE DBVOLUME 43, 48
 - DEFINE DEVCLASS 78, 87
 - DEFINE DOMAIN 107
 - DEFINE DRIVE 78, 86
 - DEFINE LIBRARY 77, 85
 - DEFINE LOGCOPY 59
 - DEFINE LOGVOLUME 43, 50

DEFINE MGMTCLASS 109
 DEFINE POLICYSET 108
 DEFINE SCHEDULE 185, 192
 DEFINE SPACETRIGGER 54
 DEFINE STGPOOL 81, 89
 DEFINE VOLUME 43, 92
 DELETE CLIENTOPT 134
 DELETE CLOPTSET 135
 DELETE DBVOLUME 60
 DELETE DEVCLASS 99
 DELETE DOMAIN 118
 DELETE FILESPACE 118, 130
 DELETE LIBRARY 99
 DELETE LOGVOLUME 61
 DELETE STGPOOL 93, 98
 DELETE VOLHISTORY 64, 188, 225
 DISABLE SESSIONS 200
 DISMOUNT VOLUME 94, 99
 ENABLE EVENTS 203
 EXPIRE INVENTORY 121, 191, 232
 EXTEND DB 60, 63
 EXTEND LOG 61, 63
 GRANT AUTHORITY 123
 HALT 200
 LABEL LIBVOLUME 45, 83, 97, 215
 MACRO 42
 MOVE DATA 101
 MOVE MEDIA 208
 QUERY ACTLOG 210, 226
 QUERY ADMIN 123, 126
 QUERY AUDITOCCUPANCY 191, 207
 QUERY CONTENT 101
 QUERY COPYGROUP 119
 QUERY DB 204, 231
 QUERY DBBACKUPTRIGGER 53
 QUERY DBVOLUME 57, 205
 QUERY DEVCLASS 62, 87
 QUERY DRIVE 207
 QUERY EVENT 194, 211
 QUERY FILESPACE 130
 QUERY LIBRARY 62
 QUERY LIBVOLUME 92, 208
 QUERY LICENSE 139
 QUERY LOG 205, 233
 QUERY LOGVOLUME 58, 206
 QUERY MEDIA 208
 QUERY NODE 129
 QUERY OCCUPANCY 102, 208
 QUERY OPTIONS 38
 QUERY PROCESS 95, 131, 201, 212
 QUERY REQUEST 91, 97, 209
 QUERY RESTORE 201, 212
 QUERY SCHEDULE 193
 QUERY SESSION 213
 QUERY STATUS 40, 52, 200, 202
 QUERY STGPOOL 82, 98, 209
 QUERY VOLHISTORY 206, 224, 225
 QUERY VOLUME 190, 209, 218, 221
 REDUCE DB 64
 REDUCE LOG 64

REGISTER ADMIN 123
 REGISTER LICENSE 137
 REGISTER NODE 127, 128
 REMOVE ADMIN 127
 REMOVE NODE 118
 RENAME STGPOOL 102
 REPLY 91, 97
 RESET BUFPOOL 232
 RESET LOGCONSUMPTION 52
 SELECT 42, 103, 133, 188
 SET ACCOUNTING 42
 SET ACTLOGRETENTION 203
 SET AUTHENTICATION 41
 SET INPWLENGTH 41
 SET INVALIDPWLIMIT 41
 SET LICENSEAUDITPERIOD 140
 SET LOGMODE 51
 SET PASSEXP 41
 SET WEBAUTHTIMEOUT 41
 UPDATE ADMIN 126
 UPDATE NODE 119, 130, 133
 UPDATE SPACETRIGGER 54, 55
 UPDATE STGPOOL 82, 98, 189, 190
 UPDATE VOLUME 220, 222
 VALIDATE POLICYSET 120

commands, client

ARCHIVE 163
 INCREMENTAL 159
 QUIT 166
 RESTORE 160
 RETRIEVE 164
 SCHEDULE 175
 SELECTIVE 159

COMMMETHOD

client option 154
 server option 144

COMPRESSALWAYS, client option 235

COMPRESSION option, UPDATE NODE command 235

configuration view, Web administrative interface 148

CONSOLE, admin parameter 199

CONTACT option, DEFINE ADMIN command 123

COPIED option, QUERY CONTENT command 101

COPY CLOPTSET command 134

D

data storage policy, management of backup data 105
 database

defining server volumes 43, 48
 extending capacity 60
 incremental and full backup 62
 managing backups 223
 MAXIMUMSIZE option 54
 mirroring volumes 56
 moving backups offsite 223
 reducing capacity 64
 removing volumes 60
 synchronization of server volumes 58
 triggering automatic backup 53
 triggering automatic expansion 54

DATEFORMAT

- client option 145
- server option 39
- DEFINE ASSOCIATION command 192, 193
- DEFINE CLIENTOPT command 134
- DEFINE CLOPTSET command 134
- DEFINE COPYGROUP command 112, 113
- DEFINE DBBACKUPTRIGGER command 53
- DEFINE DBCOPY command 58
- DEFINE DBVOLUME command 43, 48
- DEFINE DEVCLASS command 78, 87
- DEFINE DOMAIN command 107
- DEFINE DRIVE command 78, 86
- DEFINE LIBRARY command 77, 85
- DEFINE LOGCOPY command 59
- DEFINE LOGVOLUME command 43, 50
- DEFINE MGMTCLASS command 109
- DEFINE POLICYSET command 108
- DEFINE SCHEDULE command 185, 192
- DEFINE SPACETRIGGER command 54
- DEFINE STGPOOL command 81, 89
- DEFINE VOLUME command 43, 92
- DELETE CLIENTOPT command 134
- DELETE CLOPTSET command 135
- DELETE DBVOLUME command 60
- DELETE DEVCLASS command 99
- DELETE DOMAIN command 118
- DELETE FILESPACE command 118, 130
- DELETE LIBRARY command 99
- DELETE LOGVOLUME command 61
- DELETE STGPOOL command 93, 98
- DELETE VOLHISTORY command 64, 188, 225
- DELETEFILES option, ARCHIVE command 163
- DESCRIPTION option, ARCHIVE command 163
- DESTINATION option, DEFINE COPYGROUP command 111
- detail frame, Web administrative interface 148
- DEVCONFIG, server option 40
- device class
 - 8MM type 78
 - defining device class 78
 - deleting device class 99
 - DISK 79
 - DLT type 80
 - FILE type 79
 - FORMAT option 78
 - logical component of storage device 76
 - MOUNTRETENTION option 78
- devices
 - physical and logical components 76
 - random access 68, 75
 - sequential storage 69, 75
- DIRMC, client option 108, 236
- DISABLE SESSIONS command 200
- disaster recovery management
 - See DRM
- DISMOUNT VOLUME command 94, 99
- domain
 - See policy domain
- DOMAIN option, REGISTER NODE command 127
- drive

- defining for automatic library 85
- defining for manual library 77
- physical component of storage device 76
- DRM 36, 223
- dsm, backup-archive GUI program 167
- dsm.opt, client user options file 153
- dsm.sys, client system options file 153
- DSM_CONFIG, client variable 152
- DSM_DIR, client variable 152
- DSM_LOG, client variable 152
- dsmadmc, administrative client program 145
- DSMAGENT, remote client agent 152, 180
- dsmc, backup-archive command line interface client program 157
- DSMCAD, client acceptor 152, 180
- DSMCUTIL, Windows NT client service program 177
- DSMFMT, server utility 43
- DSMLABEL, server utility 45
- dsmserv, server program 42, 196
- dsmserv.opt, server options file 38
- DSMTCA, client communication agent 152

E

- ENABLE EVENTS command 203
- ENABLE3590LIBRARY, server option 40
- enterprise console
 - see Web administrative interface
- environment variables for client 152
- errpt, AIX command 227
- event monitoring 202
- event receiver 203
- event viewer, Web administrative interface 147, 148
- EXCLUDE, client option 155
- excluding files from backup processing 155
- EXPINTERVAL, server option 40, 121, 232
- EXPIRE INVENTORY command 121, 191, 232
- export-import function 35
- EXTEND DB command 60, 63
- EXTEND LOG command 61, 63

F

- file state
 - active 110, 161
 - expired 110, 161
 - inactive 110, 161
- FORMAT option, DEFINE DEVICECLASS command 78
- FORMATSIZE, volume option 43, 48, 50
- FREQUENCY option, DEFINE COPYGROUP command 111

G

- GRANT AUTHORITY command 123

H

- HALT command 200
- HELPDESK, admin ID 125, 129
- HTTPPORT
 - client option 156

server option 144

I

INACTIVE option, RESTORE command 161
INCLEXCL, client option 155
INCLUDE, client option 108, 155
include-exclude list 108, 153
including excluded files to backup processing 155
INCREMENTAL command 159
Internet Explorer, supported levels 146, 179
ITEMCOMMIT, admin parameter 42

J

Java applets 146
Java requirements
 Web administrative interface 146
 Web client 179

L

LABEL LIBVOLUME command 45, 83, 97, 215
LABELSOURCE option, LABEL LIBVOLUME command 90
LANGUAGE, client option 145
LARGECOMMBUFFERS, client option 236
LATEST option, RESTORE command 161
library
 defining automatic library to ADSM 85
 defining manual library to ADSM 77
 LIBTYPE option 77
 physical component of storage device 76
 steps to delete a library 97
LIBTYPE option, DEFINE LIBRARY command 77
licensing
 base server license
 compliance with license terms 139, 140
 enrollment certificate files 137
 evaluation try and buy 141
 migrating from evaluation to full installation 141
 NOAUDITSTORAGE server option 140
 saving in nodelock file 139
 scheduling audit 191
LOGPOOLSIZE, server option 40, 233
loop mode of command line interface 146, 158
lsdev, AIX command 77

M

MACRO command 42
macros
 loading macro into ADSM 4
 mac.admins 124
 mac.optionsets 133
 mac.policy 114
 mac.schedules 184, 188, 192
 mac.scripts 187
 mac.stgcreate 88
 mac.stgdelete 93
management class
 ARCHMC option 108

binding backup data 108
defining a management class 108
DIRMC option 108
INCLUDE option 108
MIGDESTINATION option 115
MAXIMUMSIZE option, DEFINE SPACETRIGGER command 54
MAXSCRATCH option, DEFINE STGPOOL command 80
MAXSESSION, server option 233
MAXSIZE option, DEFINE STGPOOL command 70
MIGDESTINATION option, DEFINE MGMTCLASS command 115
migration between storage pools 70, 100
mirroring of volumes 56
MIRRORWRITE, server option 40
MODE option, DEFINE COPYGROUP command 111
modified files during processing, serialization 112, 113
MOUNTRETENTION option, DEFINE DEVICECLASS command 78
MOVE DATA command 101
MOVE MEDIA command 208
MOVEBATCHSIZE, server option 40, 233
MOVESIZETHRESH, server option 40, 233

N

Netscape, supported levels 146, 179
network view, Web administrative interface 148
NEXTSTGPOOL option, DEFINE STGPOOL command 72
NOAUDITSTORAGE, server option 140
nodelock, license file 139
NODENAME, client option 154
NUMBERFORMAT
 client option 145
 server option 39

O

object view, Web administrative interface 148
operation view, Web administrative interface 148

P

PASSWORDACCESS, client option 155, 156
passwords
 automatic handling 155
 Web client 156
PATH, client variable 152
PICK option, RESTORE command 160
point-in-time restore 171
policy domain
 ARCHRET option 107
 BACKRET option 107
 defining a policy domain 107
 enforcing policy 121
 moving node to another domain 119
 retention grace period 114
 steps to remove default 118
policy set
 activating policy sets 108, 121

- defining policy sets 108
- validating policy sets 120
- POOLTYPE option, DEFINE STGPOOL command 83
- PRESERVEPATH option
 - RESTORE command 161
 - RETRIEVE command 165
- problem determination 36, 226

Q

- QUERY ACTLOG command 210, 226
- QUERY ADMIN command 123, 126
- QUERY AUDITOCCUPANCY command 191, 207
- QUERY CONTENT command 101
- QUERY COPYGROUP command 119
- QUERY DB command 204, 231
- QUERY DBBACKUPTRIGGER command 53
- QUERY DBVOLUME command 57, 205
- QUERY DEVCLASS command 62, 87
- QUERY DRIVE command 207
- QUERY EVENT command 194, 211
- QUERY FILESPACE command 130
- QUERY LIBRARY command 62
- QUERY LIBVOLUME command 92, 208
- QUERY LICENSE command 139
- QUERY LOG command 205, 233
- QUERY LOGVOLUME command 58, 206
- QUERY MEDIA command 208
- QUERY NODE command 129
- QUERY OCCUPANCY command 102, 208
- QUERY OPTIONS command 38
- QUERY PROCESS command 95, 131, 201, 212
- QUERY REQUEST command 91, 97, 209
- QUERY RESTORE command 201, 212
- QUERY SCHEDULE command 193
- QUERY SESSION command 213
- QUERY STATUS command 40, 52, 200, 202
- QUERY STGPOOL command 82, 98, 209
- QUERY VOLHISTORY command 206, 224, 225
- QUERY VOLUME command 190, 209, 218, 221
- QUIESCE option, HALT command 200
- QUIET, client option 236
- QUIT command 166

R

- RAID
- RECLAIM option
 - DEFINE STGPOOL command 71
 - UPDATE STGPOOL command 190
- RECLAIMSTGPOOL option, DEFINE STGPOOL command 71
- reclamation
 - offsite volumes 72
 - scheduling 190
 - sequential storage pool volumes 71
- recovery log
 - defining server volumes 44, 49
 - extending capacity 61
 - mirroring volumes 56
 - reducing capacity 64

- removing volumes 61
- ROLLFORWARD option 51
- saving records 51
- synchronization of volumes 59
- triggering automatic expansion 55
- redbook ADSM environment 2
- redbook support material
 - checklists
 - See checklists 5
 - creating storage pools 88
 - defining administrator 124
 - defining client option sets 133
 - defining policy structure 114
 - defining schedules 184, 188, 192
 - defining server scripts for offsite backup 187
 - deleting default storage pools 93
 - downloading from redbooks Web server 3, 9, 239, 243
 - loading macro into ADSM 4
 - planning worksheets
 - See worksheets 9
- REDUCE DB command 64
- REDUCE LOG command 64
- redundant array of independent disks
 - See RAID
- REGISTER ADMIN command 123
- REGISTER LICENSE command 137
- REGISTER NODE command 127, 128
- remote client agent, Web client 152, 180
- REMOVE ADMIN command 127
- REMOVE NODE command 118
- RENAME STGPOOL command 102
- REPLACE, client option 155
- REPLY command 91, 97
- RESET BUFPOOL command 232
- RESET LOGCONSUMPTION command 52
- restartable restore sessions 171
- RESTORE command 160
- retention grace period 114
- RETEXTRA option, DEFINE COPYGROUP command 110
- REONLY option, DEFINE COPYGROUP command 111
- RETRIEVE command 164
- RETVR option, DEFINE COPYGROUP command 113
- REVOKEREMOTEACCESS, client option
 - See Web client
- ROLLFORWARD option, SET LOGMODE command 51

S

- SCHEDMODE, client option 155
- SCHEDULE command 175
- scripts, server 35, 42, 186
- SEARCH option, QUERY ACTLOG command 226
- Secure Sockets Layer
 - see SSL
- SELECT command 42, 103, 133, 188
- SELECTIVE command 159
- SERIALIZATION option, DEFINE COPYGROUP command 112
- server

- active console mode 196
- background mode 198
- BUFPOOLSIZE option 40, 231
- COMMMETHOD option 144
- communication options 39
- database
 - See database
- DATEFORMAT option 39
- DEVCONFIG option 40
- DSMFMT utility 43
- DSMLABEL utility 45
- dsmserv program 42, 196
- dsmserv.opt, options file 38
- ENABLE3590LIBRARY option 40
- event monitoring 202
- event receiver 203
- EXPINTERVAL option 40, 121, 232
- export-import function 35
- HTTPPORT option 144
- installation 37
- LANGUAGE option 39
- licensing
 - See licensing
- loading scripts 42
- LOGPOOLSIZE option 40, 233
- MAXSESSION option 233
- MIRRORWRITE option 40
- MOVEBATCHSIZE option 40, 233
- MOVESIZETHRESH option 40, 233
- NOAUDITSTORAGE option 140
- NUMBERFORMAT option 39
- performance options 40
- recovery log
 - See recovery log
- sample scripts 42
- scripts 186
- security related settings 41
- site dependent options 39
- starting the program 195
- stopping the program 199
- storage pool
 - See storage pool
- TCPWINDOWSIZE option 40, 234
- TIMEFORMAT option 40
- TXNGROUPMAX option 40, 235
- USELARGEBUFFER option 234
- VOLUMEHISTORY option 40, 188
- volumes
 - See volumes
- SERVER_CONSOLE, admin ID 123, 196
- SET ACCOUNTING command 42
- SET ACTLOGRETENTION command 203
- SET AUTHENTICATION command 41
- SET INPWLENGTH command 41
- SET INVALIDPWLIMIT command 41
- SET LICENSEAUDITPERIOD command 140
- SET LOGMODE command 51
- SET PASSEXP command 41
- SET WEBAUTHTIMEOUT command 41
- SQL 35, 103
- SSL
 - Web administrative interface 147
- STATUS option, UPDATE VOLUME command 220
- storage pool
 - checkout of tape volumes 96
 - collection of storage pool volumes 67
 - COLLOCATE option 73
 - copy volumes 68
 - defining copy 89
 - defining disk volumes 84, 92
 - defining for device class of DISK 81
 - defining for manual library 80
 - defining primary 88
 - defining server volumes 44
 - defining volumes for automatic library 90
 - defining volumes for manual library 83
 - deleting a storage pool 98
 - devices
 - See devices
 - disk caching 73
 - labelling and checkin of tape volumes 90
 - logical component of storage device 76
 - MAXSCRATCH option 80
 - MAXSIZE option 70
 - migration 70, 100
 - moving data within storage pool 101
 - NEXTSTGPOOL option 72, 80
 - POOLTYPE option 83
 - primary volumes 68
 - RECLAIM option 71, 190
 - RECLAIMSTGPOOL option 71, 82
 - reclamation 71
 - relabelling of tape volumes 97
 - removing default 93
 - scheduling migration 189
 - single drive reclamation 71
 - storage migration hierarchy 68
 - threshold parameter for migration 100
- structured query language
 - See SQL
- SUBDIR, client option 155

T

- TAPEPROMPT, client option 155
- tapes
 - See volumes
- TCPBUFFSIZE, client option 154, 236
- TCPPORT, client option 154
- TCPSERVERADDRESS, client option 154
- TCPSERVERNAME, client option 145
- TCPWINDOWSIZE
 - client option 154, 236
 - server option 40, 234
- TIMEFORMAT
 - client option 145
 - server option 40
- tree frame, Web administrative interface 148
- trustee directory assignments, NetWare
 - backup of directory information 93
- try and buy evaluation

See licensing
TSO, administrative interface on MVS 143
TXNBYTELIMIT, client option 235, 237
TXNGROUPMAX, server option 40, 235, 237
TYPE option, DEFINE COPYGROUP command 111

U

UPDATE ADMIN command 126
UPDATE NODE command 119, 130, 133
UPDATE SPACETRIGGER command 54, 55
UPDATE STGPOOL command 82, 98, 189, 190
UPDATE VOLUME command 220, 222
USELARGEBUFFER, server option 234
USERID option, REGISTER NODE command 127

V

VALIDATE POLICYSET command 120
VERDELETED option, DEFINE COPYGROUP command 110
VEREXISTS option, DEFINE COPYGROUP command 110
VOLUMEHISTORY, server option 40, 188
volumes
 checking out from library 219
 CHECKLABEL option 94
 checkout of tapes 96
 collocation 73
 COPIED option 101
 defining disk storage pool volumes 92
 defining storage pool volumes for automatic library 90
 defining storage pool volumes for manual library 83
 FORMATSIZE option 43, 48, 50
 formatting server volumes 43
 history
 See administrative client
 labelling and checkin of tapes 90
 LABELSOURCE option 90
 listing to send offsite 218
 listing to send onsite 221
 logical component of storage device 76
 managing tapes between onsite and offsite 214
 mirroring 56
 mount request of tapes 91, 97
 moving data within storage pool 101
 moving tapes from offsite to onsite 221
 moving tapes from onsite to offsite 217
 reclamation, offsite 72, 223
 relabelling of tapes 97
 scheduling reclamation 190
 STATUS option 220

W

WAIT option, running processes in foreground 131
Web administrative interface
 accessing the Web 147
 banner frame 148
 browser requirements 146
 command line 147, 148

COMMMETHOD server option 144
configuration view 148
detail frame 148
enterprise console 146
event viewer 147, 148
HTTPPORT server option 144
Internet Explorer 146
Java applets 146
Netscape 146
network view 148
object view 148
operation view 148
software requirements 144
SSL 147
tree frame 148
WEBAUTHTIMEOUT option

Web client

 administrative authorities 180
 client acceptor 152, 180
 client access authority 180
 client functions 179
 client owner authority 180
 configuration 152
 DSMAGENT, remote client agent 152, 180
 DSMCAD, client acceptor 152, 180
 enterprise console 146
 HELPDESK admin ID 125
 HTTPPORT option 156
 Java applet 179
 PASSWORDACCESS option 156
 passwords 156
 remote client agent 152, 180
 REVOKEREMOTEACCESS client option 156
 userid 180

WEBAUTHTIMEOUT option

 See Web administrative interface

wheel of live, recommended scheduling of operations 183
WIN32 admin GUI 143
Windows NT configuration wizards 45
WINDOWS NT event viewer 228

worksheets

 administrator IDs 33
 client requirement 9
 database 26
 device configuration and volume history 28
 license requirements 33
 primary storage pool 28
 recovery log 26
 storage policy requirements 15
 tape drive configuration 29
 total ADSM disk required 29

ITSO Redbook Evaluation

Getting Started with ADSM: A Practical Implementation Guide
SG24-5416-00

Your feedback is very important to help us maintain the quality of ITSO redbooks. **Please complete this questionnaire and return it using one of the following methods:**

- Use the online evaluation form found at <http://www.redbooks.ibm.com>
- Fax this form to: USA International Access Code + 1 914 432 8264
- Send your comments in an Internet note to redbook@us.ibm.com

Which of the following best describes you?

Customer **Business Partner** **Solution Developer** **IBM employee**
 None of the above

Please rate your overall satisfaction with this book using the scale:
(1 = very good, 2 = good, 3 = average, 4 = poor, 5 = very poor)

Overall Satisfaction _____

Please answer the following questions:

Was this redbook published in time for your needs? Yes___ No___

If no, please explain:

What other redbooks would you like to see published?

Comments/Suggestions: (THANK YOU FOR YOUR FEEDBACK!)

SG24-5416-00

Printed in the U.S.A.

