

ADSTAR Distributed Storage Manager
for MVS



Administrator's Reference

Version 2

Note!

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

First Edition (July 1995)

This edition applies to Version 2 Release 1 of the ADSTAR Distributed Storage Manager (5655-119), and to any subsequent releases until otherwise indicated in new editions or technical newsletters. Make sure you are using the correct edition for the level of the product.

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

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

International Business Machines Corporation
Information Development, Dept. 61C
9000 South Rita Road
Tucson, AZ 85744-0001, U.S.A.

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

© **Copyright International Business Machines Corporation 1993, 1995. All rights reserved.**

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

Contents

Notices	ix
Trademarks	x
Preface	xi
Who Should Read This Publication	xi
What You Should Know Before Reading This Publication	xi
ADSTAR Distributed Storage Manager Publications	xi
Software Developer's Program	xii
Translations	xiii
Do You Have Comments or Suggestions?	xiii
Changes for ADSTAR Distributed Storage Manager Version 2	xv
Chapter 1. Using Syntax Diagrams	1
Chapter 2. Entering Administrative Commands	3
Command Components	3
Command Name	3
Positional Parameters	3
Keyword Parameters	3
Command Entry Rules	4
Processing Administrative Commands	5
Commands Processed in the Foreground	5
Commands Processed in the Background	5
Formatting Command Output on a Terminal	6
Using Wildcard Characters	6
Examples for Match-Any Characters	7
Examples for Match-Exactly-One Character	7
Naming Conventions for ADSM	8
Naming Volumes	8
Case-Sensitivity	8
Naming Restrictions for Objects	8
Chapter 3. Using the Command Line Interface	11
Using the Server Console Session	11
Using the Administrative Client Session	11
Using Console Mode	12
Using Batch Mode	13
Using Interactive Mode	14
Using Administrative Client Options	15
Using Continuation Characters	16
Redirecting Command Output to a File	16
Using Macros to Issue Commands	17
Writing Commands in a Macro File	18
Writing Comments in a Macro File	19
Using Substitution Variables in a Macro	19
Running the Commands in a Macro	20
Controlling Command Processing in a Macro	20
Chapter 4. Using Commands Based on Privilege Class	23

Commands Requiring System Privilege	23
Commands Requiring Policy Privilege	26
Commands Requiring Storage Privilege	29
Commands Requiring Operator Privilege	31
Commands Requiring Analyst Privilege	32
Commands Any Administrator Can Issue	33
Chapter 5. Using Administrator Commands	37
ACTIVATE POLICYSET (Activate a New Policy Set)	38
ASSIGN DEFMGMTCLASS (Assign a Default Management Class)	40
AUDIT LICENSES (Audit for Compliance with License Terms)	43
AUDIT VOLUME (Verify Database Information for a Storage Pool Volume)	45
BACKUP DB (Back Up the Database)	50
BACKUP DEVCONFIG (Create Backup Copies of Device Configuration Information)	54
BACKUP STGPOOL (Back Up Primary Storage Pool to Copy Storage Pool)	56
BACKUP VOLHISTORY (Save Sequential Volume History Information)	60
CANCEL MOUNT (Cancel a Mount Operation)	62
CANCEL PROCESS (Cancel an Administrative Process)	64
CANCEL SESSION (Cancel One or More Client Sessions)	66
COMMIT (Control Committing of Commands in a Macro)	68
COPY DOMAIN (Copy a Policy Domain)	70
COPY MGMTCLASS (Copy a Management Class)	73
COPY POLICYSET (Copy a Policy Set)	75
COPY SCHEDULE (Copy a Client or an Administrative Command Schedule)	77
Copying Client Schedules	77
Copying Administrative Command Schedules	78
DEFINE ASSOCIATION (Associate Client Nodes With a Schedule)	81
DEFINE COPYGROUP (Define a Copy Group)	83
DEFINE DBBACKUPTRIGGER (Define the Database Backup Trigger)	91
DEFINE DBCOPY (Define a Volume Copy of a Database Volume)	95
DEFINE DBVOLUME (Define a Database Volume)	97
DEFINE DEVCLASS (Define a Device Class)	99
DEFINE DOMAIN (Define a New Policy Domain)	107
DEFINE LOGCOPY (Define a Volume Copy of a Recovery Log Volume)	110
DEFINE LOGVOLUME (Define a Recovery Log Volume)	112
DEFINE MGMTCLASS (Define a Management Class)	114
DEFINE POLICYSET (Define a Policy Set)	118
DEFINE SCHEDULE (Define a Client or an Administrative Command Schedule)	121
Defining Client Schedules	121
Defining Administrative Command Schedules	127
DEFINE STGPOOL (Define a Storage Pool)	135
DEFINE VOLUME (Define a Volume in a Storage Pool)	148
DELETE ASSOCIATION (Delete Node Associations with a Schedule)	151
DELETE COPYGROUP (Delete a Copy Group)	153
DELETE DBBACKUPTRIGGER (Delete the Database Backup Trigger)	156
DELETE DBVOLUME (Delete a Database Volume)	158
DELETE DEVCLASS (Delete a Device Class)	160
DELETE DOMAIN (Delete a Policy Domain)	162
DELETE EVENT (Delete Event Records)	164
DELETE FILESPACE (Delete Client Node Data from the Server)	167
DELETE LOGVOLUME (Delete a Recovery Log Volume)	170
DELETE MGMTCLASS (Delete a Management Class)	172

DELETE POLICYSET (Delete a Policy Set)	174
DELETE SCHEDULE (Delete a Client or an Administrative Command Schedule)	176
Deleting Client Schedules	176
Deleting Administrative Command Schedules	176
DELETE STGPOOL (Delete a Storage Pool)	179
DELETE VOLHISTORY (Delete Sequential Volume History Information Collected by the Server)	181
DELETE VOLUME (Delete a Storage Pool Volume)	185
DISABLE (Temporarily Prevent Client Node Access to the Server)	188
DISMOUNT DEVICE (Dismount a Volume by Real Device Address)	190
DISMOUNT VOLUME (Dismount a Volume by Volume Name)	191
ENABLE (Resume User Activity on the Server)	192
EXPIRE INVENTORY (Manually Start Inventory Expiration Processing)	193
EXPORT ADMIN (Copy Administrator Information to Sequential Media)	195
EXPORT NODE (Copy Client Node Information to Sequential Media)	199
EXPORT POLICY (Copy Policy Information to Sequential Media)	204
EXPORT SERVER (Copy Server Information to Sequential Media)	208
EXTEND DB (Increase the Assigned Capacity of the Database)	213
EXTEND LOG (Increase the Assigned Capacity of the Recovery Log)	215
GRANT AUTHORITY (Add Administrator Authority)	217
HALT (Shut down the Server)	221
HELP (Get Help on Commands and Error Messages)	224
IMPORT ADMIN (Import Administrator Information)	227
IMPORT NODE (Import Client Node Information)	230
IMPORT POLICY (Import Policy Information)	235
IMPORT SERVER (Import the Server)	239
LOCK ADMIN (Lock Out an Administrator)	244
LOCK NODE (Lock Out a Client Node)	245
MACRO (Invoke a Macro)	246
MOVE DATA (Move Files on a Storage Pool Volume)	249
QUERY ACTLOG (Search Activity Log for Messages)	252
QUERY ADMIN (Display Information on One or More Administrators)	256
QUERY ASSOCIATION (Query Client Node Associations with a Schedule)	260
QUERY AUDIT OCCUPANCY (Query Client Node Storage Utilization)	262
QUERY CONTENT (Query the Contents of a Storage Pool Volume)	265
QUERY COPYGROUP (Query One or More Copy Groups)	270
QUERY DB (Display Information on the Database)	274
QUERY DBBACKUPTRIGGER (Query the Database Backup Trigger)	278
QUERY DBVOLUME (Display Information on One or More Database Volumes)	280
QUERY DEVCLASS (Display Information on One or More Device Classes)	283
QUERY DOMAIN (Query One or More Policy Domains)	287
QUERY EVENT (Query Scheduled and Completed Events)	290
Querying Events for Client Schedules	290
Querying Events for Administrative Command Schedules	294
QUERY FILESPACE (Query One or More File Spaces)	300
QUERY LICENSE (Display License Information)	304
QUERY LOG (Display Information on the Recovery Log)	307
QUERY LOGVOLUME (Display Information on One or More Log Volumes)	311
QUERY MGMTCLASS (Query a Management Class)	314
QUERY MOUNT (Display Information on Mounted Sequential Access Volumes)	318
QUERY NODE (Query One or More Client Nodes)	320

QUERY OCCUPANCY (Display File Space Information by Storage Pool)	324
QUERY OPTION (Query One or More Server Options)	327
QUERY POLICYSET (Query a Policy Set)	329
QUERY PROCESS (Query One or More Server Processes)	332
QUERY SCHEDULE (Query One or More Schedules)	335
QUERY SESSION (Query One or More Client Sessions)	339
QUERY STATUS (Query System Parameters)	342
QUERY STGPOOL (Query One or More Storage Pools)	347
QUERY VOLHISTORY (Display Sequential Volume History Information)	353
QUERY VOLUME (Query One or More Storage Pool Volumes)	359
QUIT (End the Interactive Mode of the Administrative Client)	365
REDUCE DB (Decrease the Assigned Capacity of the Database)	366
REDUCE LOG (Decrease the Assigned Capacity of the Recovery Log)	368
REGISTER ADMIN (Register an Administrator)	370
REGISTER LICENSE (Register a New License)	372
REGISTER NODE (Register a Client Node)	374
REMOVE ADMIN (Delete an Administrator)	377
REMOVE NODE (Delete a Client Node)	378
RENAME ADMIN (Rename an Administrator)	380
RENAME FILESPACE (Rename a Client Filespace on the Server)	382
RENAME NODE (Rename a Client Node)	384
RESET BUFPOOL (Reset the Buffer Pool Statistics for the Database)	386
RESET DBMAXUTILIZATION (Reset the Utilization Statistic for the Database)	388
RESET LOGCONSUMPTION (Reset the Cumulative Log Consumption Statistic)	389
RESET LOGMAXUTILIZATION (Reset the Utilization Statistic for the Log)	390
RESTORE STGPOOL (Restore Storage Pool Data from a Copy Pool)	391
RESTORE VOLUME (Restore Primary Volume Data from a Copy Pool)	396
REVOKE AUTHORITY (Remove Administrator Authority)	401
ROLLBACK (Rollback Uncommitted Changes in a Macro)	405
SET ACCOUNTING (Set Accounting Records On or Off)	407
SET ACTLOGRETENTION (Set the Retention Period for the Activity Log)	409
SET AUTHENTICATION (Set Password Authentication)	411
SET EVENTRETENTION (Specify the Retention Period for Event Records)	412
SET LICENSEAUDITPERIOD (Set License Audit Period)	414
SET LOGMODE (Set Mode for Saving Recovery Log Records)	416
SET MAXCMDRETRIES (Set Number of Times Scheduler Retries Commands)	419
SET MAXSCHEDESESSIONS (Set Maximum Scheduled Sessions)	421
SET PASSEXP (Set Password Expiration Date)	423
SET QUERYSCHEDPERIOD (Set Query Period for Polling Client Nodes)	424
SET RANDOMIZE (Randomly Distribute Scheduled Start Times)	426
SET REGISTRATION (Set Open or Closed Registration)	428
SET RETRYPERIOD (Set Time between Retry Attempts)	430
SET SCHEDMODES (Select a Central Scheduling Mode)	432
SET SERVERNAME (Specify the Server Name)	434
UNLOCK ADMIN (Unlock an Administrator)	435
UNLOCK NODE (Unlock a Client Node)	436
UPDATE ADMIN (Update an Administrator)	437
UPDATE COPYGROUP (Update a Copy Group)	439
UPDATE DBBACKUPTRIGGER (Update the Database Backup Trigger)	447
UPDATE DEVCLASS (Change the Attributes of a Device Class)	450
UPDATE DOMAIN (Update a Policy Domain)	458
UPDATE MGMTCLASS (Update a Management Class)	460

UPDATE NODE (Update a Client Node)	464
UPDATE POLICYSET (Update a Policy Set)	467
UPDATE SCHEDULE (Update a Client or an Administrative Command Schedule)	469
Updating Client Schedules	469
Updating Administrative Command Schedules	475
UPDATE STGPOOL (Change the Attributes of a Storage Pool)	482
UPDATE VOLHISTORY (Add the Location of a Database Dump, Database Backup, or Export Volume)	493
UPDATE VOLUME (Change the Attributes of One or More Storage Pool Volumes)	495
VALIDATE POLICYSET (Verify a Policy Set)	501
VARY (Bring a Random Access Volume Online or Offline)	504
Appendix A. Database Recovery Commands	507
DSMSERV DISPLAY DBBACKUPVOLUME (Display Database Backup Volume Information)	508
DSMSERV DISPLAY DBVOLUMES (Display Database Volume Information)	511
DSMSERV DISPLAY LOGVOLUMES (Display Recovery Log Volume Information)	513
DSMSERV RESTORE DB (Restore the Database)	515
Appendix B. Database Salvage Commands	529
DSMSERV DUMPDB (Dump the Database)	530
DSMSERV INSTALL (Install the Database and Log)	536
DSMSERV LOADDDB (Reload the Database)	538
DSMSERV AUDITDB (Verify Database Information)	544
Glossary	549
Index	565

Notices

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

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
500 Columbus Avenue
Thornwood, New York 10594, U.S.A.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
Information Enabling Requests
Dept. M13
5600 Cottle Road
San Jose, CA 95193, U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

Trademarks

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

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

The following terms are trademarks of other companies:

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

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

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

Windows is a trademark of Microsoft Corporation.

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

Preface

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

ADSM provides an automated, centrally scheduled, policy-managed backup, archive, and space management facility for file-servers and workstations.

This publication provides reference information for administrators of ADSM.

You can define and manage ADSM processes through administrator commands or a graphical user interface.

Who Should Read This Publication

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

All of the administrator commands you need to operate and maintain ADSM can be invoked either from the server console or by using an administrative client running on a workstation connected to the server.

What You Should Know Before Reading This Publication

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

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

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

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

ADSTAR Distributed Storage Manager Publications

The following table lists ADSM publications. For a complete list, see *ADSM General Information*.

The ADSM library is also available in softcopy on the *ADSM Online Product Library* CD-ROM (order number SK2T-1878).

Short Title	Publication Title	Order Number
ADSM General Information	<i>ADSTAR Distributed Storage Manager: General Information</i>	GH35-0131

Short Title	Publication Title	Order Number
ADSM Messages	<i>ADSTAR Distributed Storage Manager: Messages</i>	SH35-0133
ADSM Licensed Program Specifications	<i>ADSTAR Distributed Storage Manager for MVS: Licensed Program Specifications</i>	GH26-4038
ADSM Installing the Server and Administrative Client	<i>ADSTAR Distributed Storage Manager for MVS: Installing the Server and Administrative Client</i>	SH26-4043
ADSM Administrator's Guide	<i>ADSTAR Distributed Storage Manager for MVS: Administrator's Guide</i>	SH26-4039
ADSM Using the UNIX HSM Clients	<i>ADSTAR Distributed Storage Manager: Using the UNIX Hierarchical Storage Management Clients</i>	SH26-4030
ADSM V2 Using the Apple Macintosh Backup-Archive Client	<i>ADSTAR Distributed Storage Manager Version 2: Using the Apple Macintosh Backup-Archive Client</i>	SH26-4051
ADSM Using the UNIX Backup-Archive Clients	<i>ADSTAR Distributed Storage Manager Version 2: Using the UNIX Backup-Archive Clients</i>	SH26-4052
ADSM V2 Using the OS/2 Backup-Archive Client	<i>ADSTAR Distributed Storage Manager Version 2: Using the OS/2 Backup-Archive Client</i>	SH26-4053
ADSM V2 Using the DOS Backup-Archive Client	<i>ADSTAR Distributed Storage Manager Version 2: Using the DOS Backup-Archive Client</i>	SH26-4054
ADSM V2 Using the Novell NetWare Backup-Archive Client	<i>ADSTAR Distributed Storage Manager Version 2: Using the Novell NetWare Backup-Archive Client</i>	SH26-4055
ADSM V2 Using the Microsoft Windows Backup-Archive Clients	<i>ADSTAR Distributed Storage Manager Version 2: Using the Microsoft Windows Backup-Archive Clients</i>	SH26-4056
ADSM Using the Lotus Notes Backup Agent	<i>ADSTAR Distributed Storage Manager: Using the Lotus Notes Backup Agent</i>	SH26-4047
ADSM Installing the Clients	<i>ADSTAR Distributed Storage Manager: Installing the Clients</i>	SH26-4049
ADSM Client Reference Cards	<i>ADSTAR Distributed Storage Manager: Client Reference Cards</i>	SX26-6013

Software Developer's Program

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

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

To obtain the Software Developer's Program Information Package:

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

Translations

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

Do You Have Comments or Suggestions?

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

Changes for ADSTAR Distributed Storage Manager Version 2

The new functions for ADSM Version 2 are:

Administrative command scheduling

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

See the following changed commands for details:

```
DEFINE SCHEDULE
COPY SCHEDULE
QUERY EVENT
QUERY SCHEDULE
UPDATE SCHEDULE
```

Database backup and recovery

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

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

See the following new commands for details:

```
BACKUP DB
BACKUP DEVCONFIG
DEFINE DBBACKUPTRIGGER
DELETE DBBACKUPTRIGGER
QUERY DBBACKUPTRIGGER
RESET LOGCONSUMPTION
SET LOGMODE
UPDATE DBBACKUPTRIGGER
UPDATE VOLHISTORY
```

See also the following new offline commands for details:

```
DSMSERV DISPLAY DBBACKUPVOLUMES
DSMSERV DISPLAY DBVOLUMES
DSMSERV DISPLAY LOGVOLUMES
DSMSERV RESTORE DB
```

Two options, VOLUMEHISTORY and DEVCONFIG, can be added to the server options file to back up sequential volume history information and device configuration information to files that can be used during recovery.

The online DUMP DB command is removed.

Device class support enhancement

You can define sequential access device classes with the device type of FILE.

See the following changed commands for details:

```
DEFINE DEVCLASS
UPDATE DEVCLASS
```

Hierarchical storage management

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

See the following changed commands for details:

```
DEFINE MGMTCLASS  
UPDATE MGMTCLASS
```

Licensing

Server licensing is accomplished via options in the server options file. See the changed command REGISTER LICENSE.

Storage pool backup and recovery

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

See the following new commands for details:

```
BACKUP STGPOOL  
RESTORE STGPOOL  
RESTORE VOLUME
```

See also the following changed commands:

```
DEFINE STGPOOL  
DEFINE VOLUME  
UPDATE STGPOOL  
UPDATE VOLUME
```


Chapter 1. Using Syntax Diagrams



This section describes how to read the syntax diagrams included in the following sections for each administrator command.

To read a syntax diagram, follow the path of the line. Read from left to right and top to bottom.

- The **▶▶—** symbol indicates the beginning of a syntax diagram.
- The **—▶** symbol, at the end of a line, indicates that the syntax diagram continues onto the next line.
- The **▶—** symbol, at the beginning of a line, indicates that a syntax diagram continues from the previous line.
- The **—▶◀** symbol indicates the end of a syntax diagram.

Syntax items (for example, a keyword or variable) may be:

- Directly on the line (required)
- Above the line (default)
- Below the line (optional)

Syntax Diagram Description	Example
<p>Abbreviations: Uppercase letters denote the shortest acceptable abbreviation. If an item appears entirely in uppercase letters, it cannot be abbreviated.</p> <p>You can type the item in uppercase letters, lowercase letters, or any combination.</p> <p>In this example, you can enter CMDNA, CMDNAM, or CMDNAME in any combination of uppercase and lowercase letters.</p>	
<p>Symbols: You must code these symbols exactly as they appear in the syntax diagram.</p>	<ul style="list-style-type: none"> * Asterisk : Colon , Comma = Equal Sign - Hyphen () Parenthesis . Period
<p>Variables: Italicized lowercase items (<i>like this</i>) denote variables.</p> <p>In this example, <i>var_name</i> represents a variable you must specify when you code the CMDNAME command.</p>	

Syntax Diagram Description	Example
<p>Repetition: An arrow returning to the left means that the item can be repeated. A character within the arrow means you must separate repeated items with that character. A footnote (1) by the arrow refers to a limit that specifies how many times the item can be repeated.</p>	 <p>Note: ¹ Specify <i>repeat</i> up to 5 times.</p>
<p>Required Choices: When two or more items are in a stack and one of them is on the line, you <i>must</i> specify one item. In this example, you must choose CMDNAME=A, CMDNAME=B, or CMDNAME=C. Do not include any blanks before or after the equal sign (=).</p>	
<p>Optional Choice: When an item is below the line, the item is optional. In this example, you can choose CMDNAME=A or nothing at all. Do not include any blanks before or after the equal sign (=). When two or more items are in a stack below the line, all of them are optional. In this example, you can choose CMDNAME=A, CMDNAME=B, CMDNAME=C, or nothing at all. Do not include any blanks before or after the equal sign (=).</p>	
<p>Defaults: Defaults are above the line. The system uses the default unless you override it. You can override the default by coding an option from the stack below the line. In this example, CMDNAME=A is the default. You can also code CMDNAME=A, CMDNAME=B or CMDNAME=C. Do not include any blanks before or after the equal sign (=).</p>	
<p>Repeatable Choices: A stack of items followed by an arrow returning to the left means that you can select more than one item or, in some cases, repeat a single item. In this example, you can choose more than one VALUE, with each name delimited with a comma. Do not include any blanks before or after the equal sign (=).</p>	
<p>Syntax Fragments: Some diagrams, because of their length, must fragment the syntax. The fragment name appears between vertical bars in the diagram. The expanded fragment appears in the diagram after all the other parameters or at the bottom of the diagram. The expanded fragment is identified by a heading with the same fragment name. Commands appearing directly on the line are required. In this example, the fragment is named "Fragment"</p>	<p>Fragment:</p> 

Chapter 2. Entering Administrative Commands

This section includes the following information about entering ADSTAR Distributed Storage Manager (ADSM) commands:

- Descriptions of the components of a command
- Rules for entering administrative commands
- How ADSM processes administrative commands
- How ADSM formats command output
- Instructions for using wildcard characters
- Naming conventions for ADSM objects

Command Components

An administrator command can include one or more of these attributes:

- Command name
- Positional parameters
- Keyword parameters

The attributes of administrative commands are also discussed in Chapter 1, “Using Syntax Diagrams” on page 1.

Command Name

The first component of a command is the command name. The command name is the unique identifier for the command. It can consist of a single action word, such as HALT or ENABLE, or it can consist of an action word and an object for the action, such as DEFINE DOMAIN.

You can enter the entire command name or the abbreviation specified in the syntax diagram for the command.

Positional Parameters

Some commands have required or optional positional parameters. Positional parameters must be entered in the order indicated in the syntax diagram for a command.

For example, to copy a policy set, enter the name of the policy to which the policy set belongs, the name of the policy set you want to copy, and the name you want to assign to the copy, in that order:

```
COPY Policyset domain1 oldset newset
```

Keyword Parameters

Some commands have required or optional keyword parameters. When included in a command, a keyword parameter must follow any positional parameters in the command. A keyword parameter consists of a keyword, followed by an equal sign (=), followed by a value. Do not include any blanks before or after the equal sign (=). A keyword parameter can, for example, provide a description of an object, identify an object (such as a file or an option), or specify a setting for an option.

For example, to include a description for a new policy domain, enter:

```
DEFine D0main acctg DESCription="Accounting"
```

You can enter either the entire keyword or the abbreviation identified in the syntax diagram for the command.

Command Entry Rules

The following are general rules for entering ADSM administrative commands:

- Enter a command in any column of the input line.
- Unless otherwise noted, commands can be specified in any combination of uppercase and lowercase characters.
- If a parameter allows a list of values, enter the values immediately following the equal sign and delimit the value with a comma. For example,

```
Classes=poliCy,storagE
```

- If a description value (a string of text) for a parameter begins with a single or double quote or contains any embedded blanks or equal signs, surround the value in either single (') or double (") quotation marks. The opening and closing quotation marks must be the same type of quotation marks. For example, if the opening quotation is a single quotation mark, the closing quotation mark must also be a single quotation mark. See the following table for details on how to enter description values.

For example, to register a new client node with a description for the CONTACT parameter, enter:

```
REGister Node louie2 secret CONtact="manager"
```

The following table presents ways of entering a description value for the CONTACT parameter, using the REGISTER NODE command as an example.

The value can contain quotation marks, embedded blanks, or equal signs.

For this description	Enter this
manager	CONtact=manager
manager's	CONtact="manager's" or CONtact='manager' 's'
"manager"	CONtact= ' "manager" ' CONtact=" " "manager" " "
manager's data	CONtact="manager's data" or CONtact='manager' 's data '
manager's "data"	CONtact='manager' 's "data" '
manager=2GF	CONtact='manager=2GF'
manager reporting to 2GF	CONtact='manager reporting to 2GF' or CONtact="manager reporting to 2GF"

Processing Administrative Commands

ADSM processes administrator commands in two ways:

- Foreground
- Background

ADSM sends you a message indicating:

- Whether the command completed successfully (in the case of a foreground process)
- Whether the command has been sent to the server to be processed (in the case of a background process)
- A confirmation message, in the case of commands that affect the availability of the server or data managed by the server. (ADSM does not issue a confirmation message if a command is issued from the server console or from an administrative client with the NOCONFIRM option.)

Commands Processed in the Foreground

Foreground processing means that each command you issue must complete before you can issue another command. For example, when you enter a QUERY command, the information you request is sent to your screen. You cannot issue another command until this happens. Additional commands can be issued from different administrative clients while your command is processed.

You cannot cancel commands that are processed in the foreground. When you issue a QUERY command from the administrative client, multiple screens of output may be generated. If this occurs and additional output is not needed, the administrator can cancel the display of output to the client workstation. While canceling does not end the processing of the command, it ends any further display of command output.

Commands Processed in the Background

Background processing permits you to enter commands while the server is processing a previously issued command.

When you enter a command that creates a background process, ADSM submits a request to the server to perform the task identified in the command. The following is a list of commands that can generate background processes:

AUDIT LICENSES	EXPORT POLICY
AUDIT VOLUME	EXPORT SERVER
BACKUP DB	EXTEND LOG
BACKUP STGPOOL	IMPORT ADMIN
DEFINE DBCOPY	IMPORT NODE
DEFINE LOGCOPY	IMPORT POLICY
DELETE DBVOLUME	IMPORT SERVER
DELETE FILESPACE	MOVE DATA
DELETE LOGVOLUME	REDUCE DB
DELETE VOLUME	REDUCE LOG
EXPIRE INVENTORY	RESTORE STGPOOL
EXPORT ADMIN	RESTORE VOLUME
EXPORT NODE	VARY

Each background process is assigned a process number. You can cancel commands that generate background processes by using this process number.

Table 1 identifies the commands you use to obtain the status and process number of a background process and how to cancel a background process. If a background process is active when you cancel it, the server stops the process. Any changes that are uncommitted are rolled back, however, changes that are committed are not rolled back.

Table 1. Querying and Canceling a Process

Command	Description	Reference
CANCEL PROCESS	Cancels a background server process.	Page 64
QUERY PROCESS	Displays information about one or all background server processes.	Page 332

Formatting Command Output on a Terminal

ADSM formats command output according to your screen or window width. If the width of your screen or window is not wide enough to display the output horizontally, ADSM arranges and displays the information vertically.

Using Wildcard Characters

In some commands, such as the QUERY commands, you use wildcard characters to create a pattern-matching expression that specifies more than one object. Using wildcard characters makes it easier to tailor a command to your needs.

The wildcard characters you use depend on the operating system from which you issue commands. For example, you can use wildcard characters such as an asterisk (*) to match any (0 or more) characters or you can use a question mark (?), or a percent sign (%) to match exactly one character.

Table 2 provides references to wildcard characters for some operating systems. Use the wildcard characters appropriate for your system.

Table 2 (Page 1 of 2). Wildcard Characters by Environment

Environment	Match-Any	Match-Exactly-One
AIX/UNIX	*	?
HP	*	%
AS/400	*	%
CMS	*	%
DOS	*	?
OS/2	*	?
SunOS	*	?
TSO	*	%

Table 2 (Page 2 of 2). Wildcard Characters by Environment

Environment	Match-Any	Match-Exactly-One
Windows	*	?

For example, if you want to query all the management classes in all the policy sets in DOMAIN1 whose names begin with DEV, and your system uses an asterisk as the *match-any* character, you can enter:

```
Query MGmtclass domain1 * dev*
```

Or, if you want to query the first nine management classes in POLICYSET1 in DOMAIN1, and POLICYSET1 contains twelve management classes named MC1 through MC12, and your system uses a question mark as the match-one character, you can enter:

```
Query MGmtclass domain1 policyset1 mc?
```

ADSM displays information about management classes named MC1 through MC9.

Examples for Match-Any Characters

Table 3 shows examples that use wildcard characters to match any characters.

Table 3. Match-Any Character

Pattern	Matches	Does Not Match
ab*	ab, abb, abxxx	a, b, aa, bb
ab*rs	abrs, abtrs, abrsrs	ars, aabrs, abrss
ab*ef*rs	abefrs, abefghrs	abefr, abers

Examples for Match-Exactly-One Character

Table 4 shows examples that use wildcard characters to match exactly one character. The question mark (?) can be replaced by a percent (%) if your platform uses that character instead of ?.

Table 4. Match-Exactly-One Character

Pattern	Matches	Does Not Match
ab?	abc	ab, abab, abzzzz
ab?rs	abfrs	abrs, abllrs
ab?ef?rs	abdefjrs	abefrs, abdefrs, abefjrs
ab??rs	abcdrs, abzzrs	abrs, abjrs, abkkrs

Naming Conventions for ADSM

This section provides information about naming conventions for ADSM objects.

Naming Volumes

When you use DEFINE commands to define volumes for the database, storage pools, and the recovery log, the kind of volume name you use depends on your system:

- For MVS VSAM linear data sets—Specify a standard MVS VSAM cluster name. A VSAM cluster name consists of one or more qualifiers, which are delimited by periods. The length of each qualifier is from 1–8 alphanumeric characters. The maximum length of the cluster name is 42 characters. An example is ADSM.BFS.POOL2.VOLA.
Note: The name you specify for a volume cannot contain any embedded blanks or equal signs.
- For tape cartridges—Specify a 1–6 alphanumeric tape volume name. An example is DSMT01.

For information about standard naming conventions for MVS, refer to *MSV/ESA JCL Reference*, GC28-1654.

Case-Sensitivity

When referring to file space names, specify the file space name by using the exact capitalization that is used by ADSM. To find out the correct capitalization, use the QUERY FILESPACE command.

Naming Restrictions for Objects

ADSM restricts the number and type of characters that you can use to name ADSM objects.

Note: ADSM object names cannot contain any embedded blanks or equal signs.

Maximum Length

The following list shows the maximum length permitted for naming objects in ADSM:

Type of Name	Maximum Length
Administrators	64
Device classes	30
Management classes	30
Nodes	64
Passwords	64
Policy domains	30
Policy sets	30
Schedules	30
Server name	64
Storage pools	30

Valid Characters

The following characters are available for defining ADSM object names:

Characters

A–Z

0–9

–

.

-

+

&

Character Names

Any letter, A through Z

Any number, 0 through 9

Underscore

Period

Hyphen

Plus

Ampersand

Chapter 3. Using the Command Line Interface

This chapter provides information about:

- Using the server console session
- Using the administrative client session
- Using the administrative client options
- Using continuation characters to issue long commands
- Redirecting command output to a file
- Using macros to issue commands

Using the Server Console Session

ADSM has set up a special administrative user ID named `SERVER_CONSOLE` that allows you to administer the server from the server console after ADSM is installed. `SERVER_CONSOLE` is automatically registered as an administrator and is given system authority. This allows you to quickly and easily register and grant system privileges to other administrators as well as issue other ADSM administrative commands. (Authentication must be turned off in order for `SERVER_CONSOLE` to be used as an administrative client ID.)

After you register other system administrators, reduce the authority of `SERVER_CONSOLE` to operator privilege to restrict access to administrative functions. It is recommended that you run the server from one of your newly registered administrative clients, rather than continuing to use `SERVER_CONSOLE`.

`SERVER_CONSOLE` attributes cannot be modified; therefore, you cannot make the following changes to this user ID:

- Register or update `SERVER_CONSOLE`
- Lock or unlock `SERVER_CONSOLE` from ADSM
- Rename `SERVER_CONSOLE`
- Remove `SERVER_CONSOLE`

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

Using the Administrative Client Session

The administrative command-line client is a program that allows administrators to control and monitor the server from a remote workstation rather than from the server console.

The administrative client program can be installed on a workstation, personal computer, or terminal running in one of the following environments:

- Digital Equipment Corporation (DECstation)
- Hewlett-Packard (HP-UX)
- IBM Advanced Interactive Executive (AIX) for the RISC System/6000
- IBM Disk Operating System (DOS)
- IBM MVS Time Sharing Option (TSO)

- IBM Operating System/2 (OS/2)
- IBM VM CMS
- Microsoft Windows
- Microsoft Windows NT
- Santa Cruz Operation (SCO)
- Sun Microsystems SunOS

You can start an administrative client session using any of the following modes:

- Console
- Batch
- Interactive

Using Console Mode

Use console mode for your administrative client session when you want to monitor ADSM activities. In console mode, you can monitor basic server activities as they occur. For example, you can monitor migration processes and clients logging on to ADSM. This information is displayed on your terminal, and optionally, can be written to a file by using the OUTFILE option.

You *cannot* issue ADSM administrative commands while running in console mode.

Starting an Administrative Client in Console Mode

To start an administrative client session in console mode:

1. Enter the DSMADMC command with the CONSOLEMODE option:

```
DSMADMC / -CONSOLEMODE
```

Note: On MVS and VM operating systems, you must enter a slash (/) after the DSMADMC command, before any options.

2. Press the Enter key. The following prompt appears on your screen:

```
Enter your user id:
```

3. Enter your ADSM administrator ID, and then press the Enter key.

The following prompt appears on your screen if authentication is turned on for the server.

```
Enter your password:
```

4. Enter your ADSM administrator password, and then press the Enter key. ADSM runs in console mode until you end your session.

If you want ADSM to write all terminal output to a file, specify the OUTFILE option with a destination file name. For example, to have output written to the SAVE.OUT file, enter:

```
DSMADMC / -CONSOLEMODE -OUTFILE=save.out
```

All ADSM activity messages are written to the SAVE.OUT file.

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

```
DSMADMC / -ID=sullivan -PASSWORD=secret -CONSOLEMODE -OUTFILE=save.out
```

For more information on the ID, PASSWORD, CONSOLEMODE, and OUTFILE options, refer to “Using Administrative Client Options” on page 15.

Ending an Administrative Client Session in Console Mode

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

Environment	Break Sequence
AIX/UNIX-based clients	Ctrl+C
CMS	HX
DOS	Ctrl+C, Ctrl+Break
OS/2	Ctrl+C, Ctrl+Break
TSO	ATTN
Windows and Windows NT	Ctrl+C, Ctrl+Break

Using Batch Mode

Use batch mode for your administrative client session when you want to enter a single administrator command or a MACRO command. To start an administrative client session in batch mode, a server session must be available. Your administrative client session automatically ends when the command or macro has executed.

Starting an Administrative Client Session in Batch Mode

To start an administrative client session in batch mode, enter the DSMADMC command with your user ID and password by using the ID and PASSWORD options, respectively.

Note: On MVS and VM operating systems, you must enter a slash (/) after the DSMADMC command, before any options.

For example, to issue the ENABLE command in batch mode, enter:

```
DSMADMC / -ID=sullivan -PASSWORD=secret ENABLE
```

You can enter any of the other DSMADMC command options (*except* CONSOLEMODE) that you want to use.

If you specify a parameter with a string of text using batch mode, enclose the text in single quotation marks (' '). Double quotes are not allowed in batch mode.

In batch mode, you must enter the complete command on one line. If a command does not fit on one line and you want to enter it in batch mode, enter the command by using a macro file. If you want to use single quotation marks in batch mode, you must use them inside a macro file. For more information about using macro files, see “Using Macros to Issue Commands” on page 17.

For more information on the administrative client options, see “Using Administrative Client Options” on page 15.

Using Interactive Mode

Use interactive mode for your administrative client session when you want to enter a series of administrative commands. To start an administrative client session in interactive mode, a server session must be available. To ensure the availability of server sessions for both administrative and client node sessions, the interactive mode of the administrative client is disconnected if one or more of the following conditions is true:

- The server was stopped by using the HALT command.
- The administrative client session did not issue any commands for the length of time (in minutes) specified in the server options file.
- The administrative client session was canceled with the CANCEL SESSION command.

You can automatically restart your administrative client session by entering another command.

Starting an Administrative Client Session in Interactive Mode

To start an administrative client session in interactive mode:

1. Enter
DSMADMC
2. Press the Enter key. The following prompt appears on your screen:

```
Enter your user id:
```

3. Enter your ADSM administrator ID, and then press the Enter key. The following prompt appears on your screen if authentication is turned on for the server.

```
Enter your password:
```

4. Enter your ADSM administrator password, and then press the Enter key. The following prompt appears on your screen:

```
adsm>
```

5. Enter an administrative command each time ADSM displays the **adsm>** prompt.

When command processing is complete, the **adsm>** prompt will appear and the administrator can then perform additional work. Command processing can be stopped by either canceling the existing client session or by stopping the client program.

You can use the ID and PASSWORD options to specify your ADSM administrator ID and password if you do not want to be prompted this information. For example, enter:

```
DSMADMC / -ID=sullivan -PAssword=secret
```

In addition to the ID and PASSWORD options, you can use any of the administrative client options (except CONSOLEMODE) with the interactive mode of the

administrative client. For more information on the administrative client options, see “Using Administrative Client Options” on page 15.

You can use continuation characters in the interactive mode of the administrative client. For more information, see “Using Continuation Characters” on page 16.

Ending an Administrative Client Session in Interactive Mode

Use the QUIT command to end an administrative client session in interactive mode.

Using Administrative Client Options

In all the modes, administrative client options modify your administrative client session responses. The following administrative client options are valid from an administrative client session only:

-ID=userid

Specifies the ADSM administrator’s user ID.

-PASSWORD=password

Specifies the ADSM administrator’s password.

-NOConfirm

Specifies that you do not want ADSM to request confirmation before executing commands that affect the availability of the server or data managed by the server.

-Quiet

Specifies that you do not want ADSM to display standard output messages to your screen. However, when you specify this option, certain error messages still appear.

-CONsolemode

Specifies that you want to run ADSM in console mode. All unsolicited server console output is echoed to your terminal. You cannot enter any administrator commands while running in console mode.

-Itemcommit

Specifies that you want ADSM to commit commands inside a macro as each command is executed. For more information, see “Controlling Command Processing in a Macro” on page 20.

-OUTfile=filename

Specifies that you want ADSM to write all terminal output to a specified file. In an interactive or a console mode, output also goes to the terminal. In batch mode, output is written to the specified file only.

Besides the options listed here, you can also specify any option that is in the server options file. Each option must be preceded with a hyphen and delimited with a space. For more information on the server options file, refer to *ADSM Installing the Server and Administrative Client*.

Using Continuation Characters

You can use continuation characters in the interactive mode (loop) of the administrative client or in a macro file. Continuation characters are useful when you want to execute a command that is longer than your screen or window width. You can:

- Enter a dash (-) to indicate that you want to continue a command on the next line, or a backslash (\) at the end of the line you want to continue. For example:

```
REGister Admin pease mypasswd -  
CONtact="David, ext1234"
```

The maximum number of characters you can enter without a continuation character is 256. The maximum number of characters you can enter with continuation characters is 1500.

- If you want to continue a list of values, enter a dash or a backslash, with no preceding blank spaces, after the last comma of the list that you enter on the first line. Then, enter the remaining items in the list on the next line with no preceding blank spaces. For example:

```
STGpools=stg1,stg2,stg3,-  
stg4,stg5,stg6
```

- To continue a string of values that are enclosed in quotation marks, enter the first part of the string enclosed in quotation marks, followed by a dash or a backslash at the end of the line. Then, enter the remainder of the string on the next line enclosed in the *same* type of quotation marks. For example:

```
CONtact="DAVID PEASE, BLDG. 100, ROOM 2B, SAN JOSE, " -  
"EXT. 1234, ALTERNATE CONTACT-NORM PASS, EXT. 2345"
```

ADSM concatenates the two strings with no intervening blanks. You must use *only* this method to continue a quoted string of values across more than one line.

Redirecting Command Output to a File

You can save the output from a command by entering ADSM administrative client redirection characters at the end of the command. Redirection characters (> or >>) direct the output of a command to a file you specify instead of to your screen. Redirection is only available from the administrative client in the batch and interactive modes.

When redirecting output to a file, follow the naming conventions of the administrative client running on your operating system. For example, naming conventions for some operating systems are presented below:

- On AIX- and UNIX-based operating systems, a file name is case-sensitive and can contain as many as 64 alphanumeric characters, with unlimited use of periods (.) for delimiting the file name.
- On a DOS and Windows operating system, a file name is not case-sensitive. You can specify an 8-character, alphanumeric file name with an optional 3-character, alphanumeric extension.
- On an OS/2 operating system:
 - For file allocation table (FAT) file systems, a file name is not case-sensitive. You can specify an 8-character, alphanumeric file name with an optional 3-character, alphanumeric extension.

- For a high-performance file system (HPFS), a file name is not case-sensitive and can contain as many as 255 characters.
- On an MVS operating system, a file name is not case-sensitive. You can specify a sequential data set name, or a member of a PDS or PDSE that is from 1–44 alphanumeric characters. You can delimit the name with periods by using no more than 8 characters between each period. Each set of 8 characters must start with a letter, A through Z. For information on
For information about standard naming conventions for MVS, refer to *MSV/ESA JCL Reference*, GC28-1654.
In TSO, the normal TSO dataset prefix rules apply. In general, the user ID is added only if the file name is not enclosed in quotation marks.

A single greater than sign (>) indicates that you want to write the output to a new file or write over an existing file. For example, to write the output of a QUERY DOMAIN command to a new file, enter:

```
Query D0main acctg > dominfo.acc
```

Double greater than signs (>>) indicate that you want to append the output to the end of an existing file. For example, to append the output of a QUERY DOMAIN command to the end of an existing file, enter:

```
Query D0main acctg >> dominfo.acc
```

Redirection characters must be preceded by and followed by at least one blank space. If a file with the name you specify does *not* exist, ADSM creates a new file with that name.

You can redirect output from any command. However, the most common use for redirecting output is to save output from QUERY commands. The output from some QUERY commands can be quite lengthy, and it might be easier for you to browse the contents of a file or, in some cases, redirect the output to a file and then print the contents of the file.

Normally, ADSM formats output from QUERY commands according to your screen or window width. However, when you redirect the output of a QUERY command, it is written to the output file with a line width of about 500 characters. The output from QUERY commands can format differently in redirected output. However, only the format of the output changes. The content of the output is not modified.

Using Macros to Issue Commands

A macro is a file that contains one or more ADSM administrative commands. You can only issue macros from the administrative client in the batch or interactive modes.

The name of the macro file must follow the naming conventions of the administrative client running on your operating system. For example, the naming conventions for some operating systems are presented below:

- On AIX- and UNIX-based operating systems, a file name is case-sensitive, and can contain as many as 64 alphanumeric characters, with unlimited use of periods (.) for delimiting the file name.

- On a DOS or Windows operating system, a file name is not case-sensitive. You can specify an 8-character, alphanumeric file name with an optional 3-character, alphanumeric extension.
- On an OS/2 operating system:
 - For file allocation table (FAT) file systems, a file name is not case-sensitive. You can specify an 8-character, alphanumeric file name with an optional 3-character, alphanumeric extension.
 - For a high-performance file system (HPFS), a file name is not case sensitive and can contain as many as 255 characters.
- On an MVS operating system, a file name is not case-sensitive. You can specify a data set name that is from 1–44 alphanumeric characters. You can delimit the name with periods by using no more than 8 characters between each period. Each set of 8 characters must start with a letter, A through Z.
 In TSO, the normal TSO dataset prefix rules apply. In general, the user ID is added only if the file name is not enclosed in quotation marks.

A macro invoked from the **adsm>** prompt is called a high-level macro. Any macros invoked from within the high-level macro are called *nested* macros. You can include the MACRO command within a macro file to invoke other macros up to ten levels deep.

Creating a macro file to enter commands can be especially helpful when you want to issue commands that are used repeatedly. For example, you can create a macro file that contains a command to define a storage pool. By using variables for the keyword values in the file, you can define a storage pool without having to type all the keyword parameters.

You can also enter commands in a macro file to process related commands in a specific order. For example, you could create a macro file that contains commands to define a new schedule and to associate client nodes to that schedule. Or, you could create a macro file that contains commands to define a new policy set, validate it, and make it the active policy set for a policy domain.

In a macro that contains several commands, you can use the COMMIT and ROLLBACK commands to control command processing within the macro. For information about using those commands, see “Controlling Command Processing in a Macro” on page 20.

Writing Commands in a Macro File

When you write administrator commands in a macro, use the general rules stated in Chapter 2, “Entering Administrative Commands” on page 3. ADSM ignores any blank lines included in your macro. However, a completely blank line terminates a command that is continued (with continuation characters).

For information about how to name a macro file, see “Using Macros to Issue Commands” on page 17.

Here is an example of a macro called REG.MAC that registers and grants authority to a new administrator:

```
REGister Admin pease mypasswd -
  CONTACT='David Pease, x1234'
GRant AUTHority pease -
  CLasses=Policy,Storage -
  DOfains=domain1,domain2 -
  STGpools=stgpool1,stgpool2
```

This example uses continuation characters in the macro file to issue the commands. For more information on continuation characters, see “Using Continuation Characters” on page 16.

After you create a macro file, you can update the information it contains and use it again, or you can copy the macro file, make changes to the copy, and then run the copy.

Writing Comments in a Macro File

You can add comments to your macro file to document its purpose and functions.

To write a comment, write a slash and an asterisk (/*) to indicate the beginning of the comment, write the comment, and then write an asterisk and a slash (*/) to indicate the end of the comment. You can put a comment on a line by itself, or put it on a line that contains a command or part of a command.

For example, if you want to use a comment to identify the purpose of a macro, write:

```
/* AUTH.MAC-Register new nodes */
```

Or, if you want to write a comment to explain something about a command or part of a command, write:

```
DOmain=domain1 /*Assign node to DOMAIN1 */
```

Comments cannot be nested and cannot span lines. Every line of a comment must contain the comment delimiters.

Using Substitution Variables in a Macro

You can use substitution variables in a macro if you want to supply values for commands when you run the macro. When you use substitution variables, you can use a macro again and again, whenever you need to perform the same task for different objects or with different parameter values.

A substitution variable consists of a percent sign (%), followed by a number that indicates the number of the substitution variable. When you run the file with the MACRO command, you must replace the numbered substitution variable with a string.

For example, if you want to create a macro named AUTH.MAC to register new nodes, write:

```
/* Register new nodes */
REGister Node %1 %2 - /* userid password */
  CONTACT=%3 - /* 'name, phone number' */
  DOfain=%4 /* policy domain */
```

Then, when you run the macro, you enter the values you want to pass to the server to process the command.

If your system uses the percent sign as a wildcard character, a pattern-matching expression in a macro where the percent sign is immediately followed by a numeric digit is interpreted by the administrative client as a substitution variable.

You cannot enclose a substitution variable in quotation marks. However, a value you supply as a substitution for the variable can be a quoted string. For example, to register the node named DAVID with a password of DAVIDPW, with his name and phone number included as contact information, and assign him to the DOMAIN1 policy domain, enter:

```
MACRO auth.mac david davidpw "David Pease, x1234" domain1
```

Running the Commands in a Macro

Table 5 identifies the command you use to run a macro.

Table 5. Running a Macro

Command	Description	Reference
MACRO	Runs a specified macro file.	Page 246

You can enter the MACRO command in batch or interactive mode.

If the macro does not contain substitution variables, (such as the REG.MAC macro described in the “Writing Commands in a Macro File” on page 18), to run the macro simply enter the MACRO command with the name of the macro file. For example:

```
MACRO reg.mac
```

If the macro contains substitution variables (such as the AUTH.MAC macro described in “Using Substitution Variables in a Macro” on page 19), include the values that you want to supply after the name of the macro. Each value is delimited by a space. For example:

```
MACRO auth.mac pease mypasswd "david pease, x1234" domain1
```

If you enter fewer values than there are substitution variables in the macro, ADSM replaces the remaining variables with null strings. If you want to omit one or more values between values, enter a null string ("") for each omitted value. For example, if you omit the contact information in the previous example, you must enter:

```
MACRO auth.mac pease mypasswd "" domain1
```

Controlling Command Processing in a Macro

When you issue a MACRO command, the server processes all commands in the macro file in order, including commands contained in any nested macros. The server commits all commands in a macro after successfully completing processing for the highest-level macro. If an error occurs in any command in the macro or in any nested macro, the server terminates processing and rolls back any changes caused by all previous commands.

If you specify the ITEMCOMMIT option when you enter the DSMADMC command, the server commits each command in a macro individually, after successfully completing processing for each command. If an error occurs, the server continues processing and only rolls back changes caused by the failed command.

You can control precisely when commands are committed with the COMMIT command. If an error occurs while processing the commands in a macro, the server terminates processing of the macro and rolls back any uncommitted changes (commands that have been processed since the last COMMIT). Make sure your administrative client session is *not* running with the ITEMCOMMIT option if you want to control command processing with the COMMIT command.

Note: Commands that start background processes cannot be rolled back. For a list of commands that can start background processes, see “Commands Processed in the Background” on page 5.

You can test a macro before implementing it by using the ROLLBACK command.

You can enter the commands (except the COMMIT command) you want to issue in the macro, and enter ROLLBACK as the last command. Then, you can run the macro to verify that all the commands process successfully. Any changes to the database caused by the commands are rolled back by the ROLLBACK command you have included at the end. Remember to remove the ROLLBACK command before you make the macro available for actual use. Also, make sure your administrative client session is *not* running with the ITEMCOMMIT option if you want to control command processing with the ROLLBACK command.

If you have a series of commands that process successfully via the command line, but are unsuccessful when issued within a macro, there are probably dependencies between commands. It is possible that a command issued within a macro cannot be processed successfully until a previous command, that is issued within the same macro, is committed. Either of the following actions allow successful processing of these commands within a macro:

- Insert a COMMIT command before the command dependent on a previous command. For example, if COMMAND C is dependent upon COMMAND B, insert a COMMIT command before COMMAND C. An example of this macro is:

```
COMMAND A  
COMMAND B  
COMMIT  
COMMAND C
```

- Start the administrative client session using the ITEMCOMMIT option. This causes each command within a macro to be committed before the next command is processed.

Table 6 identifies commands you can use to commit and roll back commands in a macro.

Table 6. Committing or Rolling Back a Command in a Macro

Command	Description	Reference
COMMIT	Makes changes to the database permanent.	Page 68
ROLLBACK	Discards any uncommitted changes to the database since the last COMMIT was executed.	Page 405

Chapter 4. Using Commands Based on Privilege Class

In ADSM, there are five administrator privilege classes:

- System
- Policy
- Storage
- Operator
- Analyst

The authority granted to an administrator determines the administrative commands that the administrator can issue. After an administrator has been registered, the administrator can issue a limited set of commands, including all QUERY commands.

When you install ADSM, the server console is defined as a system administrator named SERVER_CONSOLE and is granted system privilege. For information on how to issue commands from the SERVER_CONSOLE administrative ID, refer to *ADSM Installing the Server and Administrative Client*.

The following sections describe each type of administrator privilege and the commands that can be issued by an administrator who has been granted the corresponding authority.

Commands Requiring System Privilege

A system administrator has the highest level of authority in ADSM. A system administrator can issue any administrative command and has authority to manage all policy domains and all storage pools.

Table 7 lists the commands that only a system administrator can issue.

Table 7 (Page 1 of 4). System Privilege Commands

Command	Description	Reference
AUDIT LICENSES	Checks for compliance with defined licenses.	Page 43
CANCEL PROCESS	Cancels a background server process.	Page 64
COPY DOMAIN	Creates a copy of an existing policy domain.	Page 70
COPY SCHEDULE (See note.)	Creates a copy of an existing schedule.	Page 77
DEFINE DOMAIN	Defines a policy domain that can be assigned to one or more clients.	Page 107
DEFINE SCHEDULE (See note.)	Defines a schedule for a client operation or an administrative command.	Page 121

Table 7 (Page 2 of 4). System Privilege Commands

Command	Description	Reference
DEFINE STGPOOL	Defines a storage pool as a named collection of data storage media.	Page 135
DELETE DOMAIN	Deletes a policy domain along with any policy objects in the policy domain.	Page 162
DELETE SCHEDULE (See note.)	Deletes a schedule from the database.	Page 176
DELETE STGPOOL	Deletes a storage pool from data storage.	Page 179
EXPIRE INVENTORY	Manually starts inventory expiration processing.	Page 193
EXPORT ADMIN	Copies administrative information to external media.	Page 195
EXPORT NODE	Copies client node information to external media.	Page 199
EXPORT POLICY	Copies policy information to external media.	Page 204
EXPORT SERVER	Copies all or part of the server to external media.	Page 208
GRANT AUTHORITY	Assigns one or more privilege classes to an administrator.	Page 217
IMPORT ADMIN	Restores administrative information from external media.	Page 227
IMPORT NODE	Restores client node information from external media.	Page 230
IMPORT POLICY	Restores policy information from external media.	Page 235
IMPORT SERVER	Restores all or part of the server from external media.	Page 239
LOCK ADMIN	Prevents an administrator from accessing ADSM.	Page 244
REGISTER ADMIN	Defines a new administrator without granting administrative authority.	Page 370
REGISTER LICENSE	Registers a new license with the ADSM server.	Page 372

Table 7 (Page 3 of 4). System Privilege Commands

Command	Description	Reference
REMOVE ADMIN	Removes an administrator from the list of registered administrators.	Page 377
RENAME ADMIN	Changes an ADSM administrator's name.	Page 380
REVOKE AUTHORITY	Revokes one or more privilege classes or restricts access to policy domains and storage pools.	Page 401
SET ACCOUNTING	Specifies whether accounting records are created at the end of each client session.	Page 407
SET ACTLOGRETENTION	Specifies the number of days to retain log records in the activity log.	Page 409
SET AUTHENTICATION	Specifies whether clients are prompted for a password at the beginning of a session.	Page 411
SET EVENTRETENTION	Specifies the number of days to retain event records.	Page 412
SET LICENSEAUDITPERIOD	Specifies the number of days between automatic license audits.	Page 414
SET LOGMODE	Sets the mode for saving records in the recovery log.	Page 416
SET MAXCMDRETRIES	Specifies the maximum number of retries after a failed attempt to execute a scheduled command.	Page 419
SET MAXSCHEDESESSIONS	Specifies the maximum number of client/server sessions available for processing scheduled work.	Page 421
SET PASSEXP	Specifies the number of days after which a password is expired and must be changed.	Page 423
SET QUERYSCHEDPERIOD	Specifies the frequency for clients to obtain scheduled work, in client-polling mode.	Page 424

Table 7 (Page 4 of 4). System Privilege Commands

Command	Description	Reference
SET RANDOMIZE	Specifies the randomization of start times within a window for schedules in client-polling mode.	Page 426
SET REGISTRATION	Specifies whether users can register themselves or must be registered by an administrator.	Page 428
SET RETRYPERIOD	Specifies the time between retry attempts by the client scheduler.	Page 430
SET SCHEDMODES	Specifies the central scheduling mode for the server.	Page 432
SET SERVERNAME	Specifies the name by which the server is identified.	Page 434
UNLOCK ADMIN	Enables a locked administrator to access ADSM.	Page 435
UPDATE ADMIN	Changes the password or contact information associated with any administrator.	Page 437
UPDATE SCHEDULE (See note.)	Changes the attributes of a schedule.	Page 469

Note: Indicates that this command is restricted by the authority granted to an administrator. Storage privilege is required only for administrative command schedules. Policy privilege is required for client operation schedules.

Commands Requiring Policy Privilege

A policy administrator can issue commands that relate to policy management objects such as policy domains, policy sets, management classes, copy groups, and schedules. The policy domains that a policy administrator can manage depend on the authority granted to the administrator by the system administrator.

As a policy administrator, you can have unrestricted or restricted policy privilege.

Unrestricted policy privilege permits you to issue all of the administrator commands that require policy privilege. You can issue commands that affect all existing policy domains as well as any policy domains that are defined in the future. An unrestricted policy administrator cannot define, delete, or copy policy domains.

Restricted policy privilege permits you to issue administrator commands that affect one or more policy domains for which you have been explicitly granted authority. For example, the DELETE MGMTCLASS command requires you to have policy privilege for the policy domain to which the management class belongs.

Table 8 on page 27 lists the commands that a policy administrator can issue.

Table 8 (Page 1 of 3). Policy Administrator Commands

Command	Description	Reference
ACTIVATE POLICYSET	Validates and activates a policy set within a policy domain.	Page 38
ASSIGN DEFMGMTCLASS	Assigns an existing management class as the default for a specified policy set.	Page 40
COPY MGMTCLASS	Creates a copy of an existing management class.	Page 73
COPY POLICYSET	Creates a copy of an existing policy set.	Page 75
COPY SCHEDULE (See note 2.)	Creates a copy of an existing schedule.	Page 77
DEFINE ASSOCIATION	Associates one or more clients with a schedule.	Page 81
DEFINE COPYGROUP	Defines a copy group for backup or archive processing within a specified management class.	Page 83
DEFINE MGMTCLASS	Defines a management class within the policy domain and policy set specified.	Page 114
DEFINE POLICYSET	Defines a policy set within the specified policy domain.	Page 118
DEFINE SCHEDULE (See note 2.)	Defines a schedule for a client operation or an administrative command.	Page 121
DELETE ASSOCIATION	Deletes the association between one or more clients and a schedule.	Page 151
DELETE COPYGROUP	Deletes a backup or archive copy group from a policy domain and policy set.	Page 153
DELETE EVENT (See note 1.)	Deletes event records prior to a specified date and time.	Page 164
DELETE FILESPACE	Deletes data associated with one or more client's filespaces.	Page 167

Table 8 (Page 2 of 3). Policy Administrator Commands

Command	Description	Reference
DELETE MGMTCLASS	Deletes a management class and its copy groups from a policy domain and policy set.	Page 172
DELETE POLICYSET	Deletes a policy set, including its management classes and copy groups, from a policy domain.	Page 174
DELETE SCHEDULE (See note 2.)	Deletes a schedule from the database.	Page 176
LOCK NODE	Prevents a client from accessing the server.	Page 245
REGISTER NODE	Defines a client to the server and sets options for that user.	Page 374
REMOVE NODE	Removes a client from the list of registered nodes for a specific policy domain.	Page 378
RENAME FILESPACE	Renames a client filesystem on the server.	Page 382
RENAME NODE	Changes the name for an existing user that belongs to a specific policy domain.	Page 384
UNLOCK NODE	Enables a locked user in a specific policy domain to access the server.	Page 436
UPDATE COPYGROUP	Changes one or more attributes of a copy group.	Page 439
UPDATE DOMAIN	Changes the attributes of a policy domain.	Page 458
UPDATE MGMTCLASS	Changes the attributes of a management class.	Page 460
UPDATE NODE	Changes the attributes associated with a client node.	Page 464
UPDATE POLICYSET	Changes the description of a policy set.	Page 467
UPDATE SCHEDULE (See note 2.)	Changes the attributes of a schedule.	Page 469

Table 8 (Page 3 of 3). Policy Administrator Commands

Command	Description	Reference
VALIDATE POLICYSET	Verifies and reports on conditions the administrator must consider before activating the policy set.	Page 501

Notes:

1. Indicates that this command cannot be restricted by policy domain. An administrator with unrestricted policy privilege or restricted policy privilege for a specified policy domain can issue this command.
2. Indicates that this command is restricted by the authority granted to an administrator. Storage privilege is required only for administrative command schedules. Policy privilege is required for client operation schedules.

Commands Requiring Storage Privilege

A storage administrator can issue commands that allocate and control storage resources for the server. The commands that a storage administrator can issue depend on the authority granted to the administrator by the system administrator. As a storage administrator, you can have unrestricted or restricted storage privilege.

Unrestricted storage privilege permits you to issue all of the administrator commands that require storage privilege. You can issue commands that affect all existing storage pools as well as any storage pools that are defined in the future. You can also issue commands that affect the database and the recovery log. An unrestricted storage administrator cannot define or delete storage pools.

Restricted storage privilege permits you to issue administrator commands that only affect a storage pool for which you have been granted authority. For example, the DELETE VOLUME command only affects a storage pool volume that is defined to a specific storage pool.

Table 9 lists the commands that a storage administrator can issue.

Table 9 (Page 1 of 3). Storage Administrator Commands

Command	Description	Reference
AUDIT VOLUME (See note.)	Compares database and storage pool information, and optionally, resolves any inconsistencies.	Page 45
BACKUP DB	Backs up the ADSM database to sequential access volumes.	Page 50
BACKUP DEVCONFIG	Backs up ADSM device information to a file.	Page 54
BACKUP STGPOOL	Backs up a primary storage pool to a copy storage pool.	Page 56

Table 9 (Page 2 of 3). Storage Administrator Commands

Command	Description	Reference
BACKUP VOLHISTORY	Records sequential volume history information in external files.	Page 60
DEFINE DBBACKUPTRIGGER	Defines a trigger for automatic ADSM database backups.	Page 91
DEFINE DBCOPY	Creates a volume copy of a database volume.	Page 95
DEFINE DBVOLUME	Assigns a volume to be used for database storage.	Page 97
DEFINE DEVCLASS	Defines a device class that can be used for storage pools, importing, or exporting.	Page 99
DEFINE LOGCOPY	Creates a volume copy of a recovery log volume.	Page 110
DEFINE LOGVOLUME	Defines a disk volume to the recovery log.	Page 112
DEFINE VOLUME (See note.)	Assigns a volume to be used for storage within a specified storage pool.	Page 148
DELETE DBBACKUPTRIGGER	Deletes the database backup trigger.	Page 156
DELETE DBVOLUME	Deletes a volume from the database.	Page 158
DELETE DEVCLASS	Deletes a device class name.	Page 160
DELETE LOGVOLUME	Deletes a volume from the recovery log.	Page 170
DELETE VOLHISTORY	Removes sequential volume history information from the volume history file.	Page 181
DELETE VOLUME (See note.)	Deletes a volume from a storage pool.	Page 185
EXTEND DB	Extends usable database storage within the volumes previously allocated for its use.	Page 213
EXTEND LOG	Extends usable log storage within the volumes previously allocated for its use.	Page 215

Table 9 (Page 3 of 3). Storage Administrator Commands

Command	Description	Reference
MOVE DATA (See note.)	Moves data from a specified storage pool volume to another storage pool volume.	Page 249
REDUCE DB	Attempts to free up a database volume by reducing usable storage within the volumes defined.	Page 366
REDUCE LOG	Attempts to free up a log volume by reducing usable storage within the volumes defined.	Page 368
RESTORE STGPOOL	Restores files to a primary storage pool from copy storage pools.	Page 391
RESTORE VOLUME	Restores files stored on specified volumes in a primary storage pool from copy storage pools.	Page 396
UPDATE DBBACKUPTRIGGER	Changes one or more attributes of the database backup trigger.	Page 447
UPDATE DEVCLASS	Changes the attributes of a device class.	Page 450
UPDATE STGPOOL (See note.)	Changes the attributes of a storage pool.	Page 482
UPDATE VOLHISTORY	Adds or changes location information for a volume in the volume history file.	Page 493

Note: Indicates that this command can be restricted by storage pool. An administrator with unrestricted storage privilege or restricted storage privilege for a specified storage pool can issue this command.

Commands Requiring Operator Privilege

An administrator with operator privilege can issue commands that control the immediate operation of the server and the availability of storage media.

Table 10 lists the commands that an operator can issue.

Command	Description	Reference
CANCEL SESSION	Cancels one or all active sessions with the server.	Page 66
DISABLE	Prevents new clients from accessing ADSM but permits existing sessions to continue.	Page 188
DISMOUNT DEVICE	Dismounts a sequential, removable volume by the real device address.	Page 190
DISMOUNT VOLUME	Dismounts a sequential, removable volume by the volume name.	Page 191
ENABLE	Resumes server activity following the DISABLE command.	Page 192
HALT	Shuts down the server immediately. (May be preceded by the DISABLE command.)	Page 221
UPDATE VOLUME	Updates the attributes of one or more storage pool volumes.	Page 495
VARY	Specifies whether a disk volume is available to the server for use.	Page 504

Commands Requiring Analyst Privilege

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

Table 11 lists the commands that an analyst administrator can issue.

Command	Description	Reference
RESET BUFPOOL	Resets the buffer pool statistics for the database.	Page 386
RESET DBMAXUTILIZATION	Resets the maximum utilization statistic for the database.	Page 388
RESET LOGCONSUMPTION	Resets the statistic on recovery log space used.	Page 389

Table 11 (Page 2 of 2). Analyst Administrator Commands

Command	Description	Reference
RESET LOGMAXUTILIZATION	Resets the maximum utilization statistic for the recovery log.	Page 390

Commands Any Administrator Can Issue

Table 12 lists the commands that any registered administrator can issue even if that administrator has not been granted any specific administrator privileges.

Table 12 (Page 1 of 3). Commands Issued by All Administrators

Command	Description	Reference
COMMIT	Makes changes to the database permanent.	Page 68
HELP	Displays help for administrative commands.	Page 224
MACRO	Runs a specified macro file.	Page 246
QUERY ACTLOG	Searches for and displays messages from the server activity log.	Page 252
QUERY ADMIN	Displays information about one or more ADSM administrators.	Page 256
QUERY ASSOCIATION	Displays the clients associated with one or more schedules.	Page 260
QUERY AUDITOCUPANCY	Displays the data storage utilization for a client node.	Page 262
QUERY CONTENT	Displays information about one or more files currently residing in a storage pool volume.	Page 265
QUERY COPYGROUP	Displays the attributes of an existing copy group.	Page 270
QUERY DB	Displays allocation information about the database.	Page 274
QUERY DBBACKUPTRIGGER	Displays information about the database backup trigger.	Page 278
QUERY DBVOLUME	Displays information about one or more database volumes.	Page 280

Table 12 (Page 2 of 3). Commands Issued by All Administrators

Command	Description	Reference
QUERY DEVCLASS	Displays information about one or more device classes.	Page 283
QUERY DOMAIN	Displays information about one or more policy domains.	Page 287
QUERY EVENT	Displays information about scheduled and completed events for selected clients.	Page 290
QUERY FILESPACE	Displays information about data in one or more filespaces that belong to a client.	Page 300
QUERY LICENSE	Displays information about licenses and audits.	Page 304
QUERY LOG	Displays allocation information about the recovery log.	Page 307
QUERY LOGVOLUME	Displays information about one or more log volumes.	Page 311
QUERY MGMTCLASS	Displays information about one or more management classes.	Page 314
QUERY MOUNT	Displays information about mounted sequential access media.	Page 318
QUERY NODE	Displays partial or complete information about one or more clients.	Page 320
QUERY OCCUPANCY	Displays file space information by storage pool.	Page 324
QUERY OPTION	Displays information about one or more options from the server options file.	Page 327
QUERY POLICYSET	Displays information about one or more policy sets.	Page 329
QUERY PROCESS	Displays information about one or all background server processes.	Page 332
QUERY SCHEDULE	Displays information about one or more schedules.	Page 335

Table 12 (Page 3 of 3). Commands Issued by All Administrators

Command	Description	Reference
QUERY SESSION	Displays information about all active administrator and client sessions with ADSM.	Page 339
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.	Page 342
QUERY STGPOOL	Displays information about one or more storage pools.	Page 347
QUERY VOLHISTORY	Displays sequential volume history information that has been collected by the server.	Page 353
QUERY VOLUME	Displays information about one or more storage pool volumes.	Page 359
QUIT	Ends an interactive administrative client session.	Page 365
ROLLBACK	Discards any uncommitted changes to the database since the last COMMIT was executed.	Page 405

Chapter 5. Using Administrator Commands

This chapter describes the ADSM administrator commands.

Information for each command includes:

- A description of the tasks a command performs
- The administrator privilege class required to use the command
- A syntax diagram that identifies the required and optional parameters for the command
- Descriptions of each parameter of the command
- Examples using the command
- A list of related commands

ACTIVATE POLICYSET (Activate a New Policy Set)

Use the ACTIVATE POLICYSET command to specify an existing policy set as the ACTIVE policy set for a policy domain.

The ACTIVATE POLICYSET command will fail if the following conditions exist:

- A copy group within the policy set specifies a copy storage pool as a destination.
- A management class specifies a copy pool as the destination for space-managed files.
- A default management class is not defined.

When a policy set is activated, the contents of the policy set are copied to a policy set that has the reserved name ACTIVE. Once activated, there is no real relationship between the policy set that has been activated (copied to ACTIVE) and the contents of the ACTIVE policy set. The original policy set can still be modified, but the copied definitions in the ACTIVE policy set can only be modified by activating another policy set.

Because of this copying process, you should not assume that the definitions in the policy set that was last activated are the same as those in the ACTIVE policy set; the original policy set may have been modified since activation.

You can define multiple policy sets for a policy domain, but only one policy set can be active at a time. The current active policy set is replaced by the one you specify when you issue this command.

Prior to activating a policy set, issue the VALIDATE POLICYSET command. This command verifies that a policy set is complete and valid prior to activation.

You can use the QUERY DOMAIN command to find out which policy set was last activated.

Privilege Class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the policy set belongs.

Syntax

```
▶▶—ACTivate Policyset—domainname—setname————▶▶
```

Parameters

domainname

Specifies the name of the policy domain for the policy set you want to activate. This parameter is required.

setname

Specifies the name of policy set you want to activate. This parameter is required.

Examples

Task

Activate the VACATION policy set in the EMPLOYEE_RECORDS policy domain.

Command: ACTivate POLicyset employee_records vacation

Output:

```
ANR1514I Policy set VACATION activated in policy domain
EMPLOYEE_RECORDS.
```

Related Commands

Table 13 provides references to related commands.

Table 13. Commands Related to ACTIVATE POLICYSET

Command	Description	Reference
COPY POLICYSET	Creates a copy of an existing policy set.	Page 75
DEFINE POLICYSET	Defines a policy set within the specified policy domain.	Page 118
DELETE POLICYSET	Deletes a policy set, including its management classes and copy groups, from a policy domain.	Page 174
QUERY DOMAIN	Displays information about one or more policy domains.	Page 287
QUERY POLICYSET	Displays information about one or more policy sets.	Page 329
UPDATE POLICYSET	Changes the description of a policy set.	Page 467
VALIDATE POLICYSET	Verifies and reports on conditions the administrator must consider before activating the policy set.	Page 501

ASSIGN DEFMGMTCLASS (Assign a Default Management Class)

Use the ASSIGN DEFMGMTCLASS command to specify an existing management class as the default management class for a particular policy set. You must assign a default management class for a policy set before you can activate that policy set. It is recommended that the default management class contain both an archive copy group and a backup copy group.

ADSM uses the default management class when:

- A user does not specify a management class in their include-exclude list. The server uses the default management class to manage files when they are backed up or archived.
- The management class associated with a backup copy or archived file is deleted. If the default management class does not contain a backup or archive copy group, the backup retention grace period for the policy domain is used.
- The backup copy group is deleted from a management class; backed up files that are bound to that management class are then managed with the default management class (assuming that the default management class contains a backup copy group). If the default management class does not contain a backup or archive copy group, the backup retention grace period for the policy domain is used.
- The archive copy group is deleted from a management class; archived files that are bound to that management class are then managed with the default management class (assuming that the default management class contains an archive copy group). If the default management class does not contain an archive copy group, the archive retention grace period for the policy domain is used.
- A client node is reassigned to a different policy domain, and the active policy set does not contain management classes with the same name as that to which the node's backed up files were bound. The files are bound to the default management class if the management class contains a backup copy group or an archive copy group. If the default management class does not contain these copy groups, the retention grace period is used during expiration.

Privilege Class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the policy set belongs.

Syntax

```
▶▶—ASsign DEFMGmtclass—domainname—setname—classname————▶▶
```

Parameters

domainname

Specifies the name of the policy domain to which the management class belongs. This parameter is required.

setname

Specifies the name of the policy set associated with the management class to be specified as the default management class. This parameter is required. You cannot assign a default management class to the ACTIVE policy set.

classname

Specifies the name of the management class to be specified as the default management class. This parameter is required.

Examples

Task

Assign DEFAULT1 as the default management class for policy set SUMMER in the PROG1 policy domain.

Command: ASsign DEFMGmtclass prog1 summer default1

Output:

ANR1538I Default management class set to DEFAULT1 for policy domain PROG1, SUMMER.

Related Commands

Table 14 provides references to related commands.

Table 14 (Page 1 of 2). Commands Related to ASSIGN DEFMGMTCLASS

Command	Description	Reference
ACTIVATE POLICYSET	Validates and activates a policy set within a policy domain.	Page 38
COPY MGMTCLASS	Creates a copy of an existing management class.	Page 73
DEFINE COPYGROUP	Defines a copy group for backup or archive processing within a specified management class.	Page 83
DEFINE MGMTCLASS	Defines a management class within the policy domain and policy set specified.	Page 114
DEFINE POLICYSET	Defines a policy set within the specified policy domain.	Page 118
DELETE MGMTCLASS	Deletes a management class and its copy groups from a policy domain and policy set.	Page 172
QUERY COPYGROUP	Displays the attributes of an existing copy group.	Page 270

ASSIGN DEFMGMTCLASS

Table 14 (Page 2 of 2). Commands Related to ASSIGN DEFMGMTCLASS

Command	Description	Reference
QUERY MGMTCLASS	Displays information about one or more management classes.	Page 314
QUERY POLICYSET	Displays information about one or more policy sets.	Page 329
UPDATE COPYGROUP	Changes one or more attributes of a copy group.	Page 439
UPDATE MGMTCLASS	Changes the attributes of a management class.	Page 460
VALIDATE POLICYSET	Verifies and reports on conditions the administrator must consider before activating the policy set.	Page 501

AUDIT LICENSES (Audit for Compliance with License Terms)

Use the AUDIT LICENSES command to start a background process that audits the number of registered client nodes and other supported license factors. The results of this audit are compared to the terms of the license that have been defined to the server, to determine if the current server configuration is in compliance with the license terms.

If the maximum period of 30 days has elapsed since the last license audit, the audit process will continue until complete. The ADSM administrator will not be able to cancel the background process.

After the license terms are changed via the REGISTER LICENSE command, the server does not recognize these changes until a license audit occurs. There are two ways for a license audit to occur:

- The AUDIT LICENSES command is issued from the command line
- The automatic license audit (scheduled with the SET LICENSEAUDITPERIOD command) occurs

The administrator will receive a warning message when the server does not comply with the licensing terms.

The information gathered during the audit is used to determine server compliance with license terms and is available for viewing after the audit is complete. You can view this information by issuing the QUERY LICENSE command.

The AUDIT LICENSES command also may be used after new licenses have been registered with the REGISTER LICENSE command so that the new license terms are known to the server before the next automatic license audit is run.

This command creates a background process that can be canceled with the CANCEL PROCESS command. To display information on background processes, use the QUERY PROCESS command.

Privilege Class

To issue this command, you must have system privilege.

Syntax

```
▶▶—AUDit LICenses—————▶▶
```

Parameters

None.

Examples

Task

Issue the AUDIT LICENSES command.

Command: AUDit LICenses

Output:

```

ANR2817I AUDIT LICENSES: License audit started as
process 2.

ANR2825I License audit process 2 completed successfully -
10 nodes audited.
ANR2811I Audit License completed - Server is in compliance
with license terms.
    
```

Related Commands

Table 15 provides references to related commands.

Table 15. Commands Related to AUDIT LICENSES

Command	Description	Reference
CANCEL PROCESS	Cancels a background server process.	Page 64
QUERY AUDITOCCUPANCY	Displays the data storage utilization for a client node.	Page 262
QUERY LICENSE	Displays information about licenses and audits.	Page 304
QUERY PROCESS	Displays information about one or all background server processes.	Page 332
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.	Page 342
REGISTER LICENSE	Registers a new license with the ADSM server.	Page 372
SET LICENSEAUDITPERIOD	Specifies the number of days between automatic license audits.	Page 414

AUDIT VOLUME (Verify Database Information for a Storage Pool Volume)

Use the AUDIT VOLUME command to check for any inconsistencies between database information and a storage pool volume. During the audit process, the server sends processing information to the activity log and the server console.

While an audit process is active, clients cannot restore data from the specified volume or store new data to that volume.

This command can be used to audit a volume assigned either to a primary storage pool or to a copy storage pool. If a file on a volume is detected with integrity errors, handling of the file will depend on the type of storage pool to which the volume belongs, whether the FIX option is specified on this command, and whether the file is also stored on a volume assigned to other pools.

You cannot audit a volume if the volume specified on this command is being deleted from a primary or copy storage pool.

If the AUDIT VOLUME command does not detect data-integrity errors for a file that had previously been marked as *damaged*, the state of the file is reset so that the file can be used. This provides a means for resetting the state of *damaged* files if it is determined that the errors were caused by correctable hardware problems such as a dirty tape head.

To display damaged files on a specific volume, use the QUERY CONTENT command.

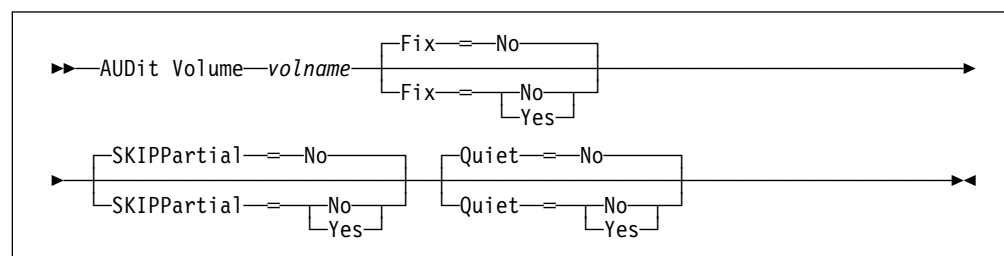
Do not use the AUDIT VOLUME command with FIX=YES if a restore process (RESTORE STGPOOL or RESTORE VOLUME) is running. The AUDIT VOLUME command could cause the restore to be incomplete.

To cancel a background process created by this command, use the CANCEL PROCESS command. To display information on background processes, use the QUERY PROCESS command.

Privilege Class

To issue this command, you must have system privilege, unrestricted storage privilege, or restricted storage privilege for the storage pool to which the volume is defined.

Syntax



Parameters

volname

Specifies the name of the volume you want to audit. This parameter is required.

Fix=fixvalue

Specifies how ADSM resolves any logical inconsistencies between the database inventory and the specified storage pool volume. This parameter is optional. The default value is NO.

The actions ADSM performs based on this parameter also depend on whether the volume is assigned to a primary or a copy storage pool.

Primary Storage Pool

For a volume in a primary storage pool, the value for the FIX parameter has the following effects.

Fix=No

ADSM reports, but does not delete, any database records that refer to files found with logical inconsistencies.

If the AUDIT VOLUME command detects a data-integrity error in a file:

- ADSM marks the file as *damaged* in the database. If a backup copy is stored in a copy storage pool, the file can be restored using the RESTORE VOLUME or RESTORE STGPOOL command.
- If the file is a cached copy, references to the file on this volume can be deleted by issuing the AUDIT VOLUME command and specifying FIX=YES.

If the AUDIT VOLUME command does not detect a data-integrity error in a file that had previously been marked as damaged, the state of the file is reset so that the file can be used. This provides a means for resetting the state of damaged files if it is determined that the errors were caused by correctable hardware problems such as a dirty tape head.

Fix=Yes

ADSM fixes any inconsistencies as they are detected.

If the AUDIT VOLUME command detects a data-integrity error in a file:

- If a backup copy is not stored in a copy storage pool, ADSM deletes all database records that refer to the file.
- If a backup copy is stored in a copy storage pool, ADSM marks the file as damaged in the database. The file can then be restored using the RESTORE VOLUME or RESTORE STGPOOL command.
- If the file is a cached copy, ADSM deletes the database records that refer to the cached file. The primary file is stored on another volume.

If the AUDIT VOLUME command does not detect a data-integrity error in a file that had previously been marked as damaged, ADSM resets the state of the file so that it can be used. This provides a means for resetting the state of damaged files if it is determined that the

errors were caused by correctable hardware problems such as a dirty tape head.

Copy Storage Pool

For a volume in a copy storage pool, the value for the FIX parameter has the following effect if ADSM detects a file with integrity errors:

Fix=No

The error is reported and the file copy is marked as *damaged* in the database.

Fix=Yes

References to the file on the audited volume are deleted. ADSM deletes any database records that point to a file on the audited volume that does not exist.

SKIPPartial=*skipvalue*

Specifies whether ADSM ignores files that span multiple storage pool volumes (*skipped files*), when performing the audit for the specified storage pool volume. This parameter is optional. The default value is NO. When performing an audit operation on a sequential access media volume, this parameter prevents additional sequential access media mounts that may be necessary to audit any *skipped* files. Possible values are:

No

Specifies that ADSM audits files that span multiple volumes when performing an audit operation on a specified volume. For sequential access volumes, multiple volumes can be mounted.

Yes

Specifies that ADSM only audits complete files that are stored on the volume to be audited. For sequential access volumes, only the specified volume is mounted. The status of any *skipped* files is unknown.

Quiet=*quietvalue*

Specifies whether ADSM sends detailed informational messages to the activity log and the server console about irretrievable files on the volume. This parameter is optional. The default value is NO. Possible values are:

No

Specifies that ADSM sends detailed informational messages and a summary report to the activity log and the server console. Each message contains the node, file space, and client name for the file.

Yes

Specifies that ADSM only sends a summary report to the activity log and the server console.

Usage Notes

Unless you specify SKIPPARTIAL=YES, ADSM attempts to process each file stored on the volume, including files that span into and out of other volumes. In order to audit files that span multiple volumes, the following conditions must be true:

- For sequential access volumes, the additional sequential access volumes must have an access mode of *read/write* or *read-only*.
- For random access volumes, the additional volumes must be online.

Examples

Task 1

Verify that the database information for storage pool volume HARVEY.ADSM.STG1 is consistent with the data stored on the volume. If any inconsistencies are found, they are not fixed.

Command: AUDit Volume HARVEY.ADSM.STG1

Output:

```
ANR2310W This command will compare all inventory references
to volume HARVEY.ADSM.STG1 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? (Y/N)
```

Type Y and press the ENTER key. A message similar to the following is displayed:

```
ANR2323I Audit Volume NOFIX process started for volume
HARVEY.ADSM.STG1 (process id 2).
```

Task 2

Verify that the database information for storage pool volume PROG2 is consistent with the data stored on the volume. ADSM fixes any inconsistencies that are found.

Command: AUDit Volume prog2 Fix=Yes

Output:

```
ANR5965I Console command:  AUDIT VOLUME JENV01 FIX=YES
ANR1199I Removable volume JENV01 is required for Audit process.
ANR5216I FILE JENV01 is expected to be mounted (R/W).
ANR2312I Audit Volume (Repair) process started for volume JENV01
        (process ID 5).
```


Related Commands

Table 16 provides references to related commands.

Table 16. Commands Related to AUDIT VOLUME

Command	Description	Reference
CANCEL PROCESS	Cancels a background server process.	Page 64
QUERY CONTENT	Displays information about one or more files currently residing in a storage pool volume.	Page 265
QUERY PROCESS	Displays information about one or all background server processes.	Page 332
QUERY VOLUME	Displays information about one or more storage pool volumes.	Page 359

BACKUP DB (Back Up the Database)

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

Copies the entire ADSM database

Incremental backup

Copies only those database pages that have been added or changed since the last time the database was backed up

If the database is damaged or lost, you can restore it by using the DSMSEV RESTORE DB command to perform one of the following types of recovery:

Roll-forward recovery

Restores the database or a single database volume to its most current state by first loading the most recent database backup series (the last full backup and all incremental backups that apply to that full backup), and then using recovery log records to apply any changes made to the database since the last backup was created. The recovery log mode must have been set to ROLLFORWARD mode continuously from the time that the last backup series was created, until the time that the database was lost or damaged. A volume history file must be available. A device configuration file must be available or must be manually created.

Point-in-time recovery

Restores the database to the point in time when a specific backup was created by using information stored in the volume history file to locate the appropriate volumes or by using the volumes you specify.

To enable use of roll-forward recovery, use the VOLUMEHISTORY and DEVCONFIG options in the server options file to automatically create volume history and device configuration files. You can also use the BACKUP VOLHISTORY and BACKUP DEVCONFIG commands. For information on using server options, see *ADSM Installing the Server and Administrative Client*.

To help you determine how much storage space a full or incremental backup will require, use the QUERY DB command. That command displays the number of changed megabytes in the database.

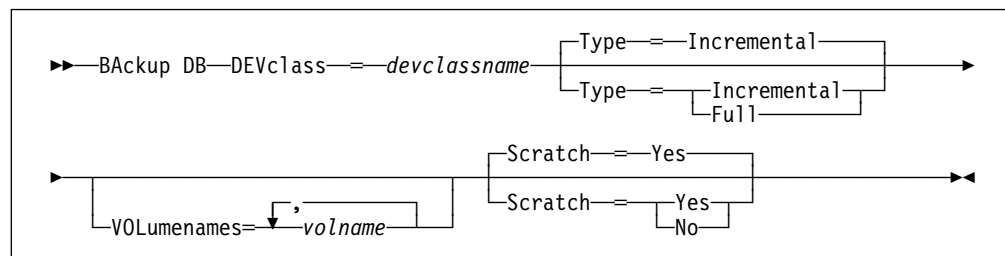
Note: If the recovery log runs out of space during a database backup and the log mode is set to ROLLFORWARD, all ADSM transactions in progress are ended and rolled back. No new transactions may begin until the backup is complete and recovery log records are deleted. After the backup is successfully completed, ADSM deletes all recovery log records except those required to restore the database to its most current state when used with the backup of the database that was just created.

This command generates a background process that can be canceled with the CANCEL PROCESS command. If a BACKUP DB background process is canceled, some of the database may have already been backed up before the cancellation. To display information on background processes, use the QUERY PROCESS command.

Privilege Class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax



Parameters

DEVclass=*devclassname*

Specifies the name of the sequential access device class to use for the backup.

Be sure that you have used the DEVCONFIG option in the server options file to specify an external file in which to store a backup copy of device class definitions. Otherwise, if your ADSM database is damaged or lost and must be restored, the definitions created by using the DEFINE DEVCLASS command will not be available and will need to be recreated manually. This parameter is required.

If the mount limit for the device class specified has already been reached (all drives are busy) when the backup runs, ADSM automatically cancels lower priority operations, such as reclamation, to make a mount point available for the backup.

Type=*typevalue*

Specifies the type of backup to run. This parameter is optional. The default value is INCREMENTAL. Possible values are:

Incremental

Specifies that you want to run an incremental backup of the ADSM database. The maximum number of incremental backups you can run between full backups is 32.

Each incremental backup, whether run automatically by ADSM based on the settings you specify with the DEFINE DBBACKUPTRIGGER command or run by issuing the BACKUP DB command, is added to the count of incremental backups run.

You can run an incremental backup whenever a full backup is not required.

Full

Specifies that you want to run a full backup of the ADSM database. A full backup is required when any of the following conditions are true:

- The database has never been backed up.

BACKUP DB

- The database has been restored using a salvage dump created with the DSMSERV DUMPDB command and has not been backed up since it was restored.
- The maximum number of incremental backups allowed between full backups has been reached.
- The log mode was changed from NORMAL to ROLLFORWARD.
- The database was extended or reduced.

Each full backup, whether run automatically by ADSM based on the settings you specify with the DEFINE DBBACKUPTRIGGER command or run by issuing the BACKUP DB command, resets the count for incremental backups to 0.

Note: Database backups (full or incremental) are not run automatically unless both the database backup trigger is set *and* the recovery log mode is set to ROLLFORWARD.

VOLUMENAMES=volname

Specifies the volumes to use for the backup. You can specify more than one volume by separating each volume name with a comma, with no intervening spaces.

This parameter is optional. However, if you specify SCRATCH=NO, you must specify a list of volumes.

Scratch=scratchvalue

Specifies whether scratch volumes can be used for the backup. This parameter is optional. The default value is YES. Possible values are:

Yes

Specifies that scratch volumes can be used.

If you include a list of volumes by using the VOLUMENAMES parameter and specify SCRATCH=YES, ADSM uses scratch volumes only if not enough space is available to store the backup data on the specified volumes.

If you do not include a list of volumes by using the VOLUMENAMES parameter, you must either specify SCRATCH=YES or omit the SCRATCH parameter (to use the default).

No

Specifies that scratch volumes cannot be used.

If you specify volumes by using the VOLUMENAMES parameter and specify SCRATCH=NO, the backup will fail if not enough space is available to store the backup data on the specified volumes.

Examples

Task

Run an incremental backup of the database, using a scratch volume. Assume a device class of FILE for the backup.

Command: BACKUP DB DEVclass=file Type=Incremental

Output:

```
ANR2281I Incremental database backup started as process 4.
```

Related Commands

Table 17 provides references to related commands.

Table 17. Commands Related to BACKUP DB

Command	Description	Reference
BACKUP DEVCONFIG	Backs up ADSM device information to a file.	Page 54
BACKUP VOLHISTORY	Records sequential volume history information in external files.	Page 60
CANCEL PROCESS	Cancels a background server process.	Page 64
DEFINE DBBACKUPTRIGGER	Defines a trigger for automatic ADSM database backups.	Page 91
DELETE DBBACKUPTRIGGER	Deletes the database backup trigger.	Page 156
DSMSERV RESTORE DB	Restores an ADSM database or database volume.	Page 515
QUERY DB	Displays allocation information about the database.	Page 274
QUERY DBBACKUPTRIGGER	Displays information about the database backup trigger.	Page 278
QUERY PROCESS	Displays information about one or all background server processes.	Page 332
SET LOGMODE	Sets the mode for saving records in the recovery log.	Page 416
UPDATE DBBACKUPTRIGGER	Changes one or more attributes of the database backup trigger.	Page 447

BACKUP DEVCONFIG (Create Backup Copies of Device Configuration Information)

Use the BACKUP DEVCONFIG command to specify one or more files in which to store a backup copy of the device class definitions created using the DEFINE DEVCLASS command.

You can run the BACKUP DEVCONFIG command at any time to ensure that device configuration information is available.


If your ADSM database is damaged or lost and must be restored using a database backup series or an offline salvage dump, device configuration information must be available before ADSM can restore the database.

You can also include the DEVCONFIG option in the server options file to specify one or more files in which to store a backup copy of device configuration information. When you use this option, the files specified are automatically updated whenever a device class is defined, updated, or deleted.

Privilege Class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax



Parameters

FileNames

Specifies the name of the file or data set in which to store a backup copy of device configuration information. You can specify more than one file or data set by separating each file name with a comma, with no intervening spaces. This parameter is optional.

If you do not specify a list of files, ADSM stores the information in all files specified with the DEVCONFIG option in the server options file.

ADSM automatically allocates the files specified if they have not been preallocated. If you choose to preallocate the files, the files must be physical sequential files.

Examples

Task

Store a backup copy of the device configuration information in a file named DEVICE.

Command: BACKUP DEVCONFig FileNames=device

Output:

```
ANR2393I BACKUP DEVCONFIG: Server device configuration
information was written to device.
```

Related Commands

Table 18 provides references to related commands.

Table 18. Commands Related to BACKUP DEVCONFIG

Command	Description	Reference
DEFINE DEVCLASS	Defines a device class that can be used for storage pools, importing, or exporting.	Page 99
UPDATE DEVCLASS	Changes the attributes of a device class.	Page 450

BACKUP STGPOOL (Back Up Primary Storage Pool to Copy Storage Pool)

Use the BACKUP STGPOOL command to create backup copies of files that reside in a primary storage pool, and store the backup copies 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 file already exists in the copy storage pool, but the copy is marked *damaged*, a new copy will be created, provided that the primary file is not also marked *damaged*.

If this command is used to back up a random-access storage pool containing cached copies of migrated files or damaged primary files, neither cached copies nor damaged primary files are backed up.

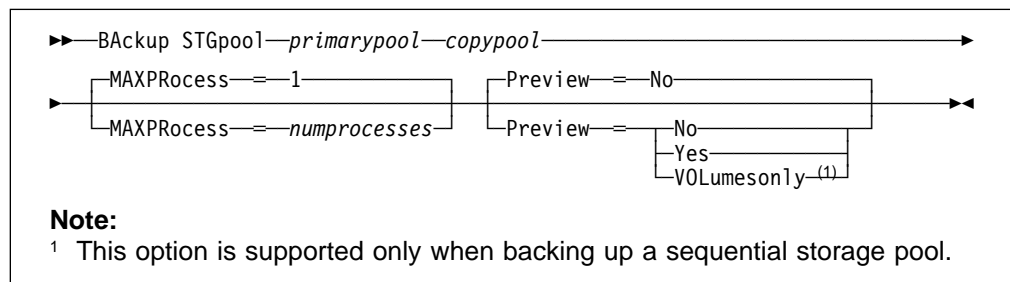
This command generates a background process that can be canceled with the CANCEL PROCESS command. If a BACKUP STGPOOL background process is canceled, some files may have already been backed up prior to the cancellation. To display information on background processes, use the QUERY PROCESS command.

Consider the hierarchy you have set up for storage pool migration when planning the backup of storage pools. If migration for a storage pool starts while that storage pool is being backed up, some files may migrate off the storage pool before they are backed up to the copy storage pool. You may want to back up storage pools that are higher in the migration hierarchy before backing up storage pools that are lower. For more information, see *ADSM Administrator's Guide*.

Privilege Class

To issue this command, you must have system privilege, unrestricted storage privilege, or restricted storage privilege for the copy storage pool in which backup copies are to be produced.

Syntax



Parameters

primarypool

Specifies the name of the primary storage pool that contains the files to be copied. This parameter is required.

copypool

Specifies the name of the copy storage pool in which the backup copies will be stored. This parameter is required.

MAXProcess=numprocesses

Specifies the maximum number of parallel processes to use for backing up files. By using multiple, parallel processes, you may be able to improve throughput for the backup operation. This parameter is optional. You can specify a value from 1–999, inclusive. The default value is 1.

When determining the number of processes to use, consider the number of mount points and drives that can be dedicated to this operation. To access a volume in a sequential access device class, ADSM uses a mount point (a logical drive), and, if the sequential device class is not FILE, a physical drive. The number of mount points and drives available depends on other ADSM and system activity, and on the mount limits of the device classes for the storage pools that are involved in the backup operation. The maximum number of mount points for a device class is specified with the mount limit parameter, which is the number of logical drives that ADSM can use at one time to access volumes in that device class.

Each process needs a mount point for copy storage pool volumes, and, if the device class is not FILE, each process also needs a drive. If the storage pool you are backing up is a sequential storage pool, each process needs an additional mount point for primary storage pool volumes, and, if the device class is not FILE, an additional drive.

For example, suppose you specify MAXPROCESS=3 to back up a primary sequential storage pool to a copy storage pool, both assigned to the same CARTRIDGE device class. Each of the three processes requires two mount points and two drives. For all three processes to run, the CARTRIDGE device class must have a mount limit of no less than 6, and there must be at least 6 mount points and 6 drives currently available.

If you specify PREVIEW=YES or PREVIEW=VOLUMESONLY, only one process is used, regardless of the value specified for the MAXPROCESS parameter. No mount points or drives are required when you specify PREVIEW=YES or PREVIEW=VOLUMESONLY.

Preview=value

Specifies whether you want to preview the results of the backup operation without actually backing up the data. This parameter is optional. The default value is NO. Possible values are:

No

Specifies that the backup operation is to be performed.

Yes

Specifies that you want to preview the results of the backup operation without actually backing up the data. This option displays the number of files and the number of bytes that will be backed up if you perform the backup. For sequential primary storage pools, this option also produces a list of primary storage pool volumes containing files that have not yet been backed up to the specified copy storage pool. These volumes must be mounted if you actually perform the backup.

Note: If you only need a list of primary storage pool volumes that must be mounted during backup, execution time for this command can be significantly reduced if you use the VOLUMESONLY option.

BACKUP STGPOOL

VOLumesonly

Specifies that you want to preview the backup operation, but only want to display a list of primary storage pool volumes containing files that have not yet been backed up to the specified copy storage pool. These volumes must be mounted if you actually perform the backup. This option is supported only for backup of sequential primary storage pools.

Examples

Task

Back up files in primary storage pool, JENSTG to copy storage pool, JENCPSTG.

Command: BBackup STGpool jenstg jencpstg

Output:

```
ANR5965I Console command: BACKUP STG JENSTG JENCPSTG
ANR1210I Backup of primary storage pool JENSTG to copy
storage pool JENCPSTG started as process 3.
ANR1228I Removable volume JENV01 is required for data
storage backup.
ANR1228I Removable volume JENV02 is required for data
storage backup.
ANR5216I FILE JENV01 is expected to be mounted (R/O).
ANR5216I FILE JENV02 is expected to be mounted (R/W).
ANR2110I BACKUP STGPOOL started as process 3.
adsm>
ANR5410I FILE volume JENV01 mounted.
ANR5410I FILE volume JENV21 mounted.
ANR5209I Dismounting volume JENV01 (read-only access).
ANR5410I FILE volume JENV02 mounted.
ANR5411I End-of-volume reached for FILE volume JENV21.
ANR5208I Dismounting volume JENV21 (updated).
ANR5410I FILE volume JENV22 mounted.
ANR1228I Removable volume JENV02 is required for data
storage backup.
ANR1212I Backup process 3 ended for storage pool JENSTG.
ANR5209I Dismounting volume JENV02 (read-onl access).
ANR5208I Dismounting volume JENV22 (updated).
ANR1214I Backup of primary storage pool JENSTG to copy
storage pool JENCPSTG has ended. Files Backed Up: 1,
Bytes Backed Up: 1613227, Unreadable
Files: 0, Unreadable Bytes: 0.
```

Related Commands

Table 19 provides references to related commands.

Table 19 (Page 1 of 2). Commands Related to BACKUP STGPOOL

Command	Description	Reference
CANCEL PROCESS	Cancels a background server process.	Page 64
QUERY PROCESS	Displays information about one or all background server processes.	Page 332

Table 19 (Page 2 of 2). Commands Related to BACKUP STGPOOL

Command	Description	Reference
QUERY STGPOOL	Displays information about one or more storage pools.	Page 347
RESTORE STGPOOL	Restores files to a primary storage pool from copy storage pools.	Page 391
RESTORE VOLUME	Restores files stored on specified volumes in a primary storage pool from copy storage pools.	Page 396

BACKUP VOLHISTORY (Save Sequential Volume History Information)

Use the BACKUP VOLHISTORY command to request that the server save sequential volume history information to one or more files.

Volume history information is collected about the following types of sequential volumes accessed by the server:

- Volumes used to back up an ADSM database (full or incremental)
- Volumes used to dump an ADSM database
- Volumes used to export administrator, node, policy, or server data
- Storage pool sequential access volumes that have been added
- Storage pool sequential access volumes that have been reused in storage pools through reclamation or MOVE DATA operations
- Storage pool sequential access volumes that have been removed by using the DELETE VOLUME command or removed during reclamation of scratch volumes

Sequential access volumes are volumes, such as tape volumes, that are accessed sequentially by the ADSM server.

To display sequential volume history information collected by the server, use the QUERY VOLHISTORY command.


To delete sequential volume history information collected by the server, use the DELETE VOLHISTORY command.

Back up volume history information to external files so that it may be used later to reload server database contents and audit affected storage pool volumes. The external files can be used by the administrator when the server cannot be started to query the database information collected for these volumes. A server options file parameter, VOLUMEHISTORY, can be used to specify one or more files that should automatically be updated whenever server sequential volume history information is changed.

Privilege Class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax



Filenames

Specifies the names of one or more files or data sets in which to store a backup copy of volume history information. Separate each file name with a comma, with no intervening spaces. This parameter is optional.

If you do not specify a file name, ADSM stores the information in all files specified with the VOLUMEHISTORY option in the server options file.

ADSM automatically allocates the files specified if they have not been preallocated. If you choose to preallocate the files, the files must be physical sequential files.

Examples

Task

Back up the volume history information in a file called VLHOST.JENV03.

Command: BAckup VOLHistory Filenames=vlhost.jenv03

Output:

```
ANR2462I BACKUP VOLHISTORY: Server sequential volume history
information was written to vlhost.jenv03.
```

Related Commands

Table 20 provides references to related commands.

Table 20. Commands Related to BACKUP VOLHISTORY

Command	Description	Reference
DELETE VOLHISTORY	Removes sequential volume history information from the volume history file.	Page 181
DELETE VOLUME	Deletes a volume from a storage pool.	Page 185
QUERY VOLHISTORY	Displays sequential volume history information that has been collected by the server.	Page 353
UPDATE VOLHISTORY	Adds or changes location information for a volume in the volume history file.	Page 493

CANCEL MOUNT

CANCEL MOUNT (Cancel a Mount Operation)

Use the CANCEL MOUNT command to cancel one or more pending tape mount requests. The mount is canceled by specifying the volume name. Any processes or sessions waiting on the mount request fail.

You can use the QUERY ACTLOG command to see which process or session failed as a result of the CANCEL MOUNT command.

Privilege Class

To issue this command, you must have system privilege or operator privilege.

Syntax

```
▶▶—CANce1—M0unt—volname—▶▶  
                          SCRTCH
```

Parameters

volname

Specifies the name of the volume for which the mount request should be canceled.

SCRTCH

Specifies that all pending scratch mount requests are to be canceled.

Examples

Task

From an administrator client, cancel the mount request for volume ST0077.

Command CANce1 M0unt st0077

Output

```
ANR4293I Cancelling mount for volume ST0077.
```

Related Commands

Table 21 provides references to related commands.

Table 21. Commands Related to CANCEL MOUNT

Command	Description	Reference
QUERY ACTLOG	Searches for and displays messages from the server activity log.	Page 252
QUERY MOUNT	Displays information about mounted sequential access media.	Page 318
QUERY PROCESS	Displays information about one or all background server processes.	Page 332

CANCEL PROCESS

CANCEL PROCESS (Cancel an Administrative Process)

Use the CANCEL PROCESS command to cancel a background process that has been started as the result of an administrative command, or from an internal server process (like migration).

To use this command, specify the process number of the process that you want to cancel. If you do not know the process number, obtain it by using the QUERY PROCESS command.

The following commands generate background processes:

AUDIT LICENSES	EXPORT SERVER
AUDIT VOLUME	EXTEND LOG
BACKUP DB	IMPORT ADMIN
BACKUP STGPOOL	IMPORT NODE
DEFINE DBCOPY	IMPORT POLICY
DEFINE LOGCOPY	IMPORT SERVER
DELETE DBVOLUME	MOVE DATA
DELETE FILESPACE	REDUCE DB
DELETE LOGVOLUME	REDUCE LOG
DELETE VOLUME	RESTORE STGPOOL
EXPIRE INVENTORY	RESTORE VOLUME
EXPORT ADMIN	VARY
EXPORT NODE	
EXPORT POLICY	

The following internal server operations generate background server processes:

- Inventory expiration
- Migration
- Reclamation

If the process you want to cancel is currently waiting for a sequential access volume to be mounted, the mount request is automatically cancelled.

Privilege Class

To issue this command you must have system privilege.

Syntax

```
▶▶—CANce1 PRocess—processnum————▶▶
```

Parameters

processnum

Specifies the number of the background process you want to cancel. This parameter is required.

Usage Notes

- If the process you want to cancel is currently waiting for a sequential access media volume to be mounted, the mount request is automatically canceled.
- If you cancel an internal server process (like migration), that process can restart immediately. For example, if you cancel a migration process, then another migration process for the same storage pool can start if additional data is stored in the storage pool (forcing migration to begin).
- To permanently cancel a migration process, you can modify the attributes of the storage pool (NEXTSTGPOOL="" or HIGHMIG=100).
- To permanently cancel a reclamation process, you can modify an attribute of the storage pool (RECLAIM=100).

Examples

Task

Cancel background process number 3.

Command: CANce1 PRocess 3

Output:

```
ANR0940I Cancel request accepted for process 3.
```

Related Commands

Table 22 provides a reference to a related command.

Table 22. Command Related to CANCEL PROCESS

Command	Description	Reference
QUERY PROCESS	Displays information about one or all background server processes.	Page 332

CANCEL SESSION (Cancel One or More Client Sessions)

Use the CANCEL SESSION command to cancel an existing administrative or client node session. You can use this command to force an administrative or client node session off of ADSM. To prevent any future client node sessions from beginning, use the DISABLE command.

When you issue this command, any outstanding mount requests related to this session are automatically canceled.

If you cancel an administrative or client node session that is in the *idle wait* (IdleW) state, the client session automatically reconnects the next time that an action is performed.

To prevent a specific administrator or client node from accessing ADSM, use the LOCK ADMIN or LOCK NODE command. If the client is currently accessing the server, the client session must end or be canceled before that node can be locked out with the LOCK NODE command.


You must cancel administrative client sessions by session number. Use the QUERY SESSION command to display information on all sessions that can be canceled, including session numbers.

When you cancel a client node session, the client node must start a new session to resume ADSM activities from the point of interruption. The results of any incomplete processing at the point of interruption are rolled back and not committed to the database. Processing that is complete prior to the point of interruption, however, is committed and is not rolled back.

Privilege Class

To issue this command, you must have system privilege or operator privilege.

Syntax



Parameters

sessionnum

Specifies the number of the administrative or client node session that you want to cancel. You must specify a session number or ALL.

ALL

Specifies that all client node sessions are canceled. You must specify a session number or ALL.

You cannot use this parameter to cancel administrative client sessions.

Usage Notes

If the session you want to cancel is currently waiting for a sequential access media volume to be mounted (in the MediaW state), the mount request is automatically canceled.

Examples

Task 1

Cancel the client node session with NODEP (session 3).

Command: CANcel SEssion 3

Output:

```
ANR0490I Canceling session 3.
```

Task 2

Cancel the administrative session with MARIE (session 5).

Command: CANcel SEssion 5

Output:

```
ANR0490I Canceling session 5.
```

Related Commands

Table 23 provides references to related commands.

Table 23. Commands Related to CANCEL SESSION

Command	Description	Reference
LOCK ADMIN	Prevents an administrator from accessing ADSM.	Page 244
LOCK NODE	Prevents a client from accessing the server.	Page 245
QUERY SESSION	Displays information about all active administrator and client sessions with ADSM.	Page 339

COMMIT (Control Committing of Commands in a Macro)

Use the COMMIT command in a macro to permanently update the database with all changes that result from the successful completion of any commands in the macro processed up to that point. This command does not generate a message when issued from the console mode of the administrative client.

You can control precisely when commands are committed with the COMMIT command. If an error occurs while processing the commands in a macro, the server terminates processing of the macro and rolls back any uncommitted changes (commands that have been processed since the last COMMIT).


Make sure that your administrative client session is not running with the ITEMCOMMIT option if you want to control command processing with this command. After a command is committed, it cannot be rolled back.

For more information on controlling command processing in a macro, see "Controlling Command Processing in a Macro" on page 20.

Privilege Class

Any administrator can issue this command.

Syntax



Parameters

None.

Examples

Task

From the interactive mode of the administrative client, register and grant authority to new administrators by using a macro named REG.ADM. Changes are committed after each administrator is registered and has been granted authority.

Macro Contents:

```
/* REG.ADM-register policy admin & grant authority*/  
REGister Admin sara hobby  
GRant AUTHority sara CLasses=Policy  
COMMIT /* Commits changes */  
REGister Admin ken plane  
GRant AUTHority ken CLasses=Policy  
COMMIT /* Commits changes */
```

Command: MACRO reg.adm

Output:

```
ANS5101I Server command: 'REGister Admin sara hobby
ANR2068I Administrator SARA registered.

ANS5101I Server command: 'GRant AUTHority sara
Classes=policy'
ANR2077I Unrestricted policy privilege granted to
administrator SARA.

ANS5101I Server command: 'COMMIT'

ANS5101I Server command: 'REGister Admin ken plane
ANR2068I Administrator KEN registered.

ANS5101I Server command: 'GRant AUTHority ken Classes=policy'
ANR2077I Unrestricted policy privilege granted to
administrator KEN.

ANS5101I Server command: 'COMMIT'
```

Related Commands

Table 24 provides references to related commands.

Table 24. Commands Related to COMMIT

Command	Description	Reference
MACRO	Runs a specified macro file.	Page 246
ROLLBACK	Discards any uncommitted changes to the database since the last COMMIT was executed.	Page 405

COPY DOMAIN (Copy a Policy Domain)

Use the COPY DOMAIN command to create a copy of an existing policy domain.

When you use this command, ADSM copies the following information:

- Policy domain description
- Policy sets in the policy domain (including the ACTIVE policy set, if defined)
- Management classes in each policy set (including the default management class, if assigned)
- Copy groups in each management class

Use the UPDATE DOMAIN command to modify an existing policy domain. DEFINE, DELETE, and UPDATE commands are available for modifying management classes and copy groups.

Privilege Class

To issue this command you must have system privilege.

Syntax

```
►►—COPy D0main—domainname1—domainname2—◄◄
```

Parameters

domainname1

Specifies the name of the existing policy domain you want to copy. This parameter is required.

domainname2

Specifies the name of the new policy domain. This parameter is required. The maximum length of this name is 30 characters.

Examples

Task 1

Copy the existing policy domain PROG1 to PROG2.

Command: COPy D0main prog1 prog2

Output:

```
ANR1503I Policy domain PROG1 copied to domain PROG2.
```

Related Commands

Table 25 provides references to related commands.

Table 25 (Page 1 of 2). Commands Related to COPY DOMAIN

Command	Description	Reference
ACTIVATE POLICYSET	Validates and activates a policy set within a policy domain.	Page 38
COPY MGMTCLASS	Creates a copy of an existing management class.	Page 73
DEFINE COPYGROUP	Defines a copy group for backup or archive processing within a specified management class.	Page 83
DEFINE DOMAIN	Defines a policy domain that can be assigned to one or more clients.	Page 107
DEFINE MGMTCLASS	Defines a management class within the policy domain and policy set specified.	Page 114
DEFINE POLICYSET	Defines a policy set within the specified policy domain.	Page 118
DELETE COPYGROUP	Deletes a backup or archive copy group from a policy domain and policy set.	Page 153
DELETE DOMAIN	Deletes a policy domain along with any policy objects in the policy domain.	Page 162
DELETE MGMTCLASS	Deletes a management class and its copy groups from a policy domain and policy set.	Page 172
QUERY COPYGROUP	Displays the attributes of an existing copy group.	Page 270
QUERY DOMAIN	Displays information about one or more policy domains.	Page 287
QUERY MGMTCLASS	Displays information about one or more management classes.	Page 314
QUERY POLICYSET	Displays information about one or more policy sets.	Page 329

COPY DOMAIN

Table 25 (Page 2 of 2). Commands Related to COPY DOMAIN

Command	Description	Reference
REGISTER NODE	Defines a client to the server and sets options for that user.	Page 374
UPDATE COPYGROUP	Changes one or more attributes of a copy group.	Page 439
UPDATE DOMAIN	Changes the attributes of a policy domain.	Page 458
UPDATE MGMTCLASS	Changes the attributes of a management class.	Page 460
UPDATE POLICYSET	Changes the description of a policy set.	Page 467
VALIDATE POLICYSET	Verifies and reports on conditions the administrator must consider before activating the policy set.	Page 501

COPY MGMTCLASS (Copy a Management Class)

Use the COPY MGMTCLASS command to create a copy of an existing management class within the same policy set.

When you use this command, ADSM copies the following information:

- Management class description
- Copy groups defined to the management class
- The following space management attributes, if they are specified in the management class:
 - SPACEMGTECHNIQUE
 - AUTOMIGNONUSE
 - MIGREQUIRESBKUP
 - MIGDESTINATION

After you copy a management class, you can use the UPDATE MGMTCLASS command to change the management class description.

You cannot copy management classes across policy domains or policy sets.

Privilege Class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the new management class belongs.

Syntax

```
►►—COPY Mgmtclass—domainname—setname—classname1—classname2—◄◄
```

Parameters

domainname

Specifies the name of the policy domain to which the management class belongs. This parameter is required.

setname

Specifies the name of the policy set to which the management class belongs. This parameter is required.

classname1

Specifies the name of the existing management class you want to copy. This parameter is required.

classname2

Specifies the name of the new management class. This parameter is required. The maximum length of this name is 30 characters.

Examples

Task

Copy the existing management class ACTIVEFILES to FILEHISTORY for policy set VACATION in the EMPLOYEE_RECORDS policy domain.

COPY MGMTCLASS

Command: COpy Mgmtclass employee_records vacation activefiles
filehistory

Output:

```
ANR1523I Management class ACTIVEFILES copied to  
class FILEHISTORY in policy domain EMPLOYEE_RECORDS,  
set VACATION.
```

Related Commands

Table 26 provides references to related commands.

Table 26. Commands Related to COPY MGMTCLASS

Command	Description	Reference
DEFINE COPYGROUP	Defines a copy group for backup or archive processing within a specified management class.	Page 83
DELETE MGMTCLASS	Deletes a management class and its copy groups from a policy domain and policy set.	Page 172
QUERY COPYGROUP	Displays the attributes of an existing copy group.	Page 270
QUERY MGMTCLASS	Displays information about one or more management classes.	Page 314
QUERY POLICYSET	Displays information about one or more policy sets.	Page 329
UPDATE COPYGROUP	Changes one or more attributes of a copy group.	Page 439
UPDATE MGMTCLASS	Changes the attributes of a management class.	Page 460

COPY POLICYSET (Copy a Policy Set)

Use the COPY POLICYSET command to create a copy of any existing policy set (including the ACTIVE policy set) within the same policy domain.

When you use this command, ADSM copies the following information:

- Policy set description
- Management classes in the policy set (including the default management class, if assigned)
- Copy groups in each management class

A policy set cannot be copied across policy domains.

To change the description of the new policy set, use the UPDATE POLICYSET command.

Privilege Class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the new policy set belongs.

Syntax

```
►►—COPY POLICYSET—domainname—setname1—setname2—◄◄
```

Parameters

domainname

Specifies the name of the policy domain to which the policy set belongs. This parameter is required.

setname1

Specifies the name of the existing policy set you want to copy. This parameter is required.

setname2

Specifies the name of the new policy set. This parameter is required. The maximum length of this name is 30 characters.

Examples

Task

Copy the existing policy set VACATION to HOLIDAY in the EMPLOYEE_RECORDS policy domain.

COPY POLICYSET

Command: COPY POLICYSET employee_records vacation holiday

Output:

```
ANR1513I Policy set VACATION copied to set HOLIDAY in
policy domain EMPLOYEE_RECORDS.
```

Related Commands

Table 27 provides references to related commands.

Table 27. Commands Related to COPY POLICYSET

Command	Description	Reference
ACTIVATE POLICYSET	Validates and activates a policy set within a policy domain.	Page 38
COPY MGMTCLASS	Creates a copy of an existing management class.	Page 73
DEFINE MGMTCLASS	Defines a management class within the policy domain and policy set specified.	Page 114
DELETE POLICYSET	Deletes a policy set, including its management classes and copy groups, from a policy domain.	Page 174
QUERY POLICYSET	Displays information about one or more policy sets.	Page 329
UPDATE POLICYSET	Changes the description of a policy set.	Page 467
VALIDATE POLICYSET	Verifies and reports on conditions the administrator must consider before activating the policy set.	Page 501

COPY SCHEDULE (Copy a Client or an Administrative Command Schedule)

Use the COPY SCHEDULE command to create a copy of an existing schedule. The COPY SCHEDULE command takes two forms, depending on whether the schedule applies to client operations or administrative commands. The syntax and parameters for each form are defined separately.

Copying Client Schedules

Use the COPY SCHEDULE command to create a copy of an existing schedule for client operations. You can copy a schedule within a policy domain or from one policy domain to another existing policy domain.

When you use this command, ADSM copies the following information:

- Description of the schedule
- All parameter values in the original schedule

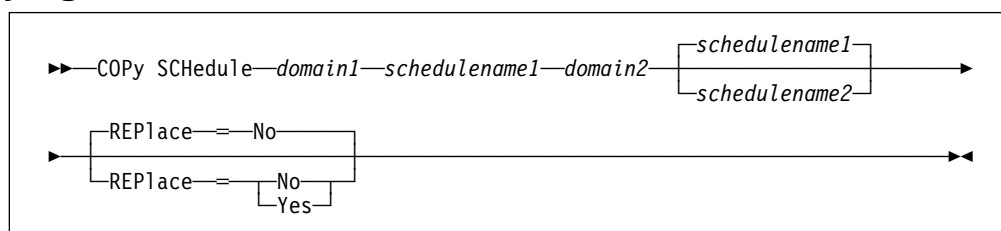
The last modification value for the new schedule is set to the time that you executed this command. Any clients associated with the original schedule are *not* associated with the new schedule. An administrator with policy privilege must use the DEFINE ASSOCIATION command to associate clients to the new schedule.

To make changes to an existing schedule, use the UPDATE SCHEDULE command.

Privilege Class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which you are copying the schedule.

Syntax for Copying Client Schedules



Parameters for Copying Client Schedules

domain1

Specifies the name of the policy domain that contains the schedule you want to copy. This parameter is required.

schedulename1

Specifies the name of the schedule you want to copy. This parameter is required.

domain2

Specifies the name of an existing policy domain to which you want to copy the new schedule. This parameter is required. You must have policy privi-

Parameters for Copying Administrative Command Schedules

schedulename1

Specifies the name of the schedule you want to copy. This parameter is required.

schedulename2

Specifies the name of the new schedule. This parameter is required. The maximum length of the name is 30 characters.

If the schedule name is already defined, you must specify REPLACE=YES, or the command fails.

Type=Administrative

Specifies this is a scheduled administrative command. This parameter is required.

REPlace=replacevalue

Specifies whether an existing schedule can be replaced. This parameter is optional. The default value is NO. Possible values are:

No

Specifies that an existing schedule cannot be replaced.

Yes

Specifies that an existing schedule can be replaced.

Examples

Task 1

Copy the WINTER schedule that belongs to policy domain DOMAIN1 to the DOMAIN2 policy domain and name the new schedule WINTERCOPY. If there is already a schedule with this name defined in the DOMAIN2 policy domain, do not replace it.

Command: COPy SCHedule domain1 winter domain2 wintercopy

Output:

```
ANR2503I Schedule WINTER in domain DOMAIN1 copied to
schedule WINTERCOPY in domain DOMAIN2.
```

Task 2

Copy the DAILYBACKUP schedule that is defined to policy domain OS2_DOMAIN1 to the OS2_DOMAIN2 policy domain. If there is already a DAILYBACKUP schedule defined to the OS2_DOMAIN2 policy domain, it is replaced.

Command: COPy SCHedule os2_domain1 dailybackup os2_domain2 REPlace=Yes

Output:

```
ANR2503I Schedule DAILYBACKUP in domain OS2_DOMAIN1
copied to schedule DAILYBACKUP in domain OS2_DOMAIN2.
```

COPY SCHEDULE

Task 3

Copy existing administrative schedule BACKUP_ARCHIVEPOOL to a new schedule named BCKSCHED.

Command: COPY SCHEDULE backup_archivepool bcksched Type=Administrative

Output:

```
ANR2513I Schedule BACKUP_ARCHIVEPOOL copied to schedule BCKSCHED.
```

Related Commands

Table 28 provides references to related commands.

Table 28. Commands Related to COPY SCHEDULE

Command	Description	Reference
DEFINE ASSOCIATION	Associates one or more clients with a schedule.	Page 81
DEFINE SCHEDULE	Defines a schedule for a client operation or an administrative command.	Page 121
DELETE SCHEDULE	Deletes a schedule from the database.	Page 176
QUERY SCHEDULE	Displays information about one or more schedules.	Page 335
UPDATE SCHEDULE	Changes the attributes of a schedule.	Page 469

DEFINE ASSOCIATION (Associate Client Nodes With a Schedule)

Use the DEFINE ASSOCIATION command to associate one or more clients with a schedule. Client nodes that are associated with a schedule initiate backup and archive functions according to that schedule.

Before associating a client to a schedule, the following conditions must be true:

- The specified schedule exists
- The client is registered to the policy domain to which the schedule applies

Any clients already associated with a specified schedule are not affected when you associate new clients with that schedule.


A client remains associated with a schedule until you use the DELETE ASSOCIATION command to remove it from the list of associated clients or until the schedule is deleted by using the DELETE SCHEDULE command.

In addition, any clients that you subsequently define to ADSM are *not* automatically associated with a schedule. You must use the DEFINE ASSOCIATION command to associate any newly registered clients with the appropriate schedules.

Privilege Class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the schedule belongs.

Syntax



Parameters

domainname

Specifies the name of the policy domain to which the schedule belongs. This parameter is required.

schedulename

Specifies the name of the schedule that you want to associate with one or more clients. This parameter is required.

nodename

Specifies the name of the client node to be associated with the specified schedule. This parameter is required. You can specify a list of clients that you want to associate with the specified schedule. The items in the list are separated by commas, with no intervening spaces. You can use a pattern-matching expression to specify a name. All matching clients are associated with the specified schedule. If a client is listed, but is already associated with the specified schedule or is not assigned to the domain to which the schedule belongs, the command has no effect for that client.

DEFINE ASSOCIATION

Examples

Task 1

Associate clients with the WEEKLY_BACKUP schedule. The associated clients have names that begin with SMITH or JOHN, and are assigned to the DOMAIN1 policy domain.

Command: DEFine ASSOCiation domain1 weekly_backup smith*,john*

Output:

```
ANR2510I Node JOHN04 associated with schedule WEEKLY_BACKUP
in policy domain DOMAIN1.
ANR2510I Node JOHNNY associated with schedule WEEKLY_BACKUP
in policy domain DOMAIN1.
ANR2510I Node SMITH associated with schedule WEEKLY_BACKUP
in policy domain DOMAIN1.
```

Task 2

Associate the clients JOE, TOM, and LARRY with the WINTER schedule, assigned to the DOMAIN1 policy domain. However, the client JOE is already associated with the WINTER SCHEDULE.

Command: DEFine ASSOCiation domain1 winter joe,tom,larry

Output:

```
ANR2510I Node TOM associated with schedule WINTER in policy
domain DOMAIN1.
ANR2510I Node LARRY associated with schedule WINTER in policy
domain DOMAIN1.
```

Related Commands

Table 29 provides references to related commands.

Table 29. Commands Related to DEFINE ASSOCIATION

Command	Description	Reference
DEFINE SCHEDULE	Defines a schedule for a client operation or an administrative command.	Page 121
DELETE ASSOCIATION	Deletes the association between one or more clients and a schedule.	Page 151
DELETE SCHEDULE	Deletes a schedule from the database.	Page 176
QUERY ASSOCIATION	Displays the clients associated with one or more schedules.	Page 260
REGISTER NODE	Defines a client to the server and sets options for that user.	Page 374

DEFINE COPYGROUP (Define a Copy Group)

Use the DEFINE COPYGROUP command to define a new backup or archive copy group within a specified policy domain, policy set, and management class. This command also defines the destination storage pool and backup or archive attributes for the copy group. When you define a copy group, it must be named STANDARD.

The DEFINE COPYGROUP will fail if a copy storage pool name is specified as a destination.

A *backup copy group* contains attributes that ADSM uses during the backup process to determine the following:

- Whether a file is only backed up when the file has changed since the last time it was backed up
- Whether a file that is in use is backed up
- How many days must elapse before a file can again be backed up
- How to process a file that is in use during backup

A backup copy group also contains attributes that ADSM uses to manage the backup versions of your files. These attributes control the following:

- Where the server stores backup versions of client node files and directories
- How many backup versions the server keeps of client node files and directories
- How long the server keeps backup versions of client node files and directories

An *archive copy group* contains attributes that ADSM uses during the archive process to determine the following:

- Whether a file that is in use is archived
- Where the server stores archived copies of client node files
- How long the server keeps archived copies of client node files

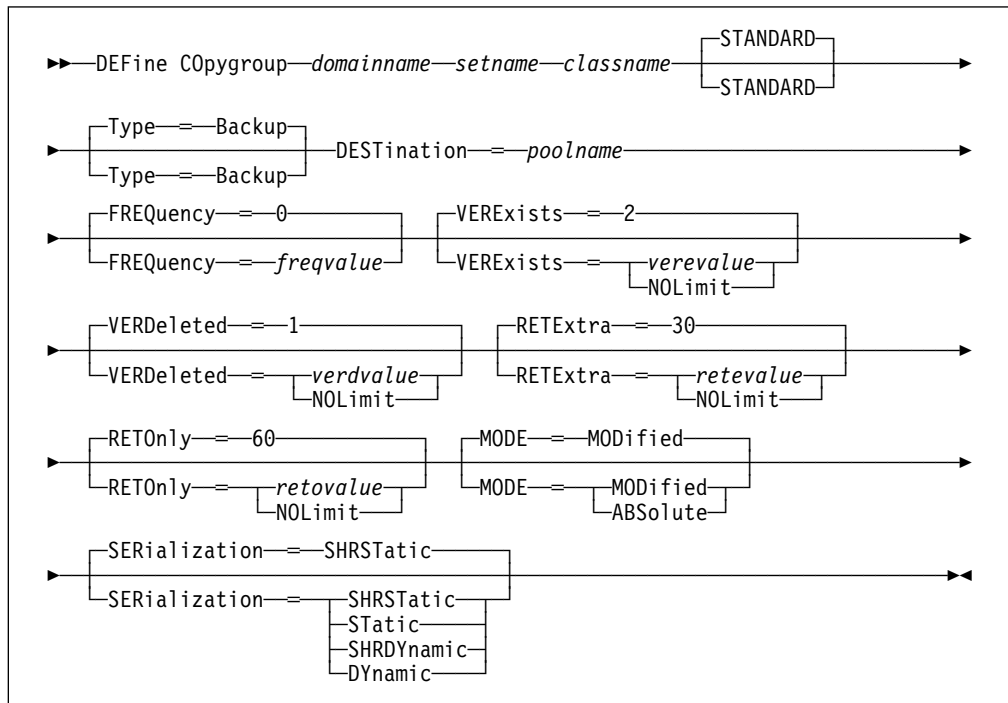
You can define one backup and one archive copy group for each management class. If you plan to have the client nodes in your installation perform backups, it is recommended that the default management class contain a backup copy group.

The DEFINE COPYGROUP command takes two forms, depending upon whether you are defining a backup copy group or an archive copy group. The syntax for each form is defined separately below.

Privilege Class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the copy group belongs.

Syntax for Defining a Backup Copy Group



Parameters for Defining a Backup Copy Group

domainname

Specifies the name of the policy domain for which you are defining the copy group. This parameter is required.

setname

Specifies the name of the policy set for which you are defining the copy group. This parameter is required.

You cannot define a copy group for a management class that belongs to the ACTIVE policy set.

classname

Specifies the name of the management class for which you are defining the copy group. This parameter is required.

STANDARD

Specifies the name of the copy group as STANDARD. The name of the copy group must be STANDARD. This parameter is optional. The default value is STANDARD.

Type=Backup

Specifies that you want to define a backup copy group. The default value is BACKUP.

DESTINATION=poolname

Specifies the name of the primary storage pool in which backup data is to be stored. This parameter is required. You can only specify one storage pool as the destination for a copy group.

FREquency=*freqvalue*

Specifies the minimum interval, in days, between successive backups. This parameter is optional. ADSM does not perform an incremental backup unless the specified number of days has elapsed since the last backup. The FREQUENCY value is only used during incremental backup. This value is ignored during selective backup. You can specify a value from 0–9999, inclusive. The default value is 0.

VERExists=*vervalue*

Specifies the maximum number of backup versions to retain for files that are currently on the client file system. This parameter is optional.

If a backup causes the limit to be exceeded, the oldest backup version that exists is deleted from the storage pool. The default value is 2. Possible values are:

vervalue

Specifies the number of backup versions to retain for files that are currently on the client file system. You can specify a value from 1–9999, inclusive. The default value is 2.

NOLimit

Specifies that you want ADSM to retain all backup versions.

VERDeleted=*verdvalue*

Specifies the maximum number of backup versions to retain for files that are deleted from the client file system. This parameter is optional. This parameter is ignored if the file still resides at the client node.

If the client file is deleted, the next backup causes the oldest versions in excess of this number to be deleted from the backup storage pool. The expiration date for the remaining versions is determined by the values specified for the RETEXTRA and RETONLY parameters. The default value is 1. Possible values are:

verdvalue

Specifies the number of backup versions to retain for files that are deleted from the client file system. You can specify a value from 0–9999, inclusive.

NOLimit

Specifies that you want ADSM to retain all backup versions for files that are deleted from the client file system.

RETExtra=*retevalue*

Specifies the retention time, in days, for all but the most recent backup version. The value of this parameter determines which versions are deleted during inventory expiration processing. This parameter is optional. The default value is 30 days. Possible values are:

retevalue

Specifies the number of days to retain all backup versions other than the most recent version. You can specify a value from 0–9999, inclusive.

NOLimit

Specifies that you want ADSM to retain all backup versions, other than the most recent version, indefinitely (that is, inventory expiration processing does not affect the backup versions). The most recent version of

DEFINE COPYGROUP

a file that exists on the client workstation is retained indefinitely or until the user deletes the file on the workstation. If the user deletes the file on the workstation the file is rolled off by using the value specified for the VERDELETED parameter.

If the NOLIMIT value is specified, extra backup versions are deleted during backup processing under control of the VEREXISTS or VERDELETED parameters.

RETO*Only=retovalue*

Specifies the retention time, in days, for the last backup version of a file that has been deleted from the client's file system. This parameter is optional. The default value is 60. Possible values are:

retovalue

Specifies the number of days to retain the last remaining inactive version of a file. You can specify a value from 0–9999, inclusive.

NOLimit

Specifies that the last remaining inactive version of a file be kept indefinitely.

MODE*=mode*

Specifies whether a file should be backed up based on changes made to the file since the last time it was backed up. This parameter is optional. The MODE value is only used for incremental backup and is ignored during selective backup. The default value is MODIFIED. Possible values are:

MODified

Specifies that the file is backed up only if it has changed since the last backup.

ABSolute

Specifies that the file is backed up regardless of whether it has been modified.

SERIALIZATION*=serialvalue*

Specifies how files or directories are handled if they are modified during backup processing and what ADSM should do if a modification occurs. This parameter is optional. The default value is SHRSTATIC.

Warning: Be careful about choosing a management class that contains a copy group that specifies SERIALIZATION=SHRDYNAMIC or SERIALIZATION=DYNAMIC.

If you choose a management class that allows a file or directory to be backed up while modifications are occurring, the backup version stored on the server might be a fuzzy backup. A fuzzy backup is a backup version that does not accurately reflect what is currently in the file or directory. The fuzzy backup contains some, but not all, modifications. If that is not acceptable, be sure to choose a management class that allows ADSM to create a backup version only if the file or directory is not being modified.

If a file is restored that contains a fuzzy backup, the file may not be usable, depending on the file's application.

Possible values are:

SHRStatic

Specifies that a file or directory must not be modified during backup. ADSM attempts to perform a backup as many as four times, depending on the value specified for CHANGINGRETRIES in the client user options file. If the file or directory is modified during each backup attempt, ADSM does not back it up.

Static

Specifies that a file or directory must not be modified during backup. If a file or directory is modified during the backup attempt, it is not backed up.

Platforms that do not support the STATIC option default to SHRSTATIC. The STATIC option is not supported on the DOS platform.

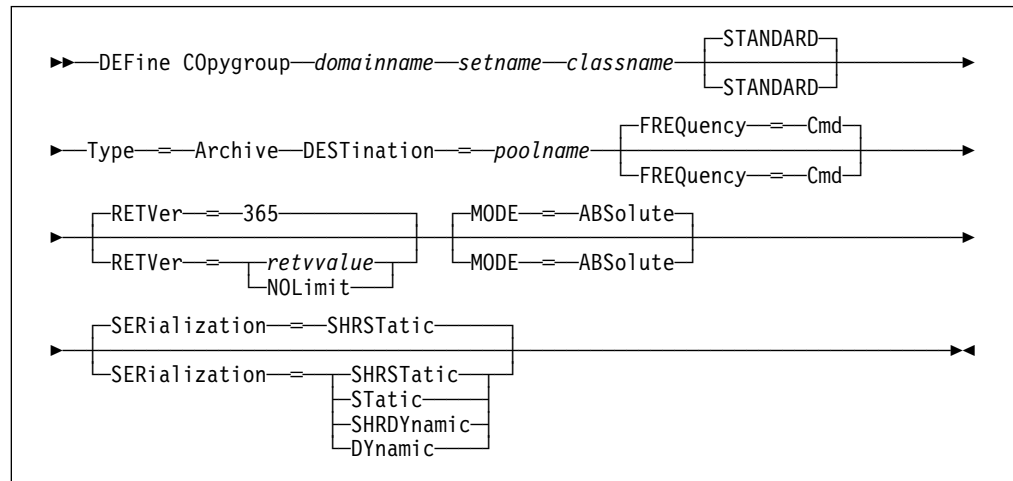
SHRDynamic

Specifies that if the file or directory is modified during a backup attempt, ADSM backs up the file or directory during the last attempt even though the file or directory is being modified. ADSM attempts to perform a backup as many as four times, depending on the value specified for CHANGINGRETRIES in the client user options file.

Dynamic

Specifies that a file or directory is backed up on the first attempt by ADSM, regardless of whether the file or directory is being modified during backup processing.

Syntax for Defining an Archive Copy Group



Parameters for Defining an Archive Copy Group

domainname

Specifies the name of the policy domain for which you are defining the copy group. This parameter is required.

setname

Specifies the name of the policy set for which you are defining the copy group. This parameter is required.

DEFINE COPYGROUP

classname

Specifies the name of the management class for which you are defining the copy group. This parameter is required.

STANDARD

Specifies the name of the copy group. The name of the copy group must be STANDARD. The default value is STANDARD.

Type=Archive

Specifies that you want to define an archive copy group. This parameter is required.

DESTination=poolname

Specifies the name of the primary storage pool where the archive files are to reside. This parameter is required. You can only specify one storage pool as the destination for a copy group.

FREQuency=Cmd

Specifies the value for the copy frequency. The copy frequency must be CMD. The parameter is optional. The default value is CMD.

RETVer=retvvalue

Specifies the number of days to keep an archive copy bound to this management class. The value of this parameter determines whether an archive copy is deleted during inventory expiration processing. This parameter is optional. The default value is 365. Possible values are:

retvvalue

Specifies the number of days to retain an archive copy. You can specify a value from 0–9999, inclusive.

NOLimit

Specifies that an archive copy is maintained indefinitely (that is, inventory expiration processing does not affect the archive copy).

MODE=ABSolute

Specifies that a file is archived. The file must be archived as ABSOLUTE. This parameter is optional. The default value is ABSOLUTE.

SERialization=serialvalue

Specifies how files are handled if they are modified during archive processing and what ADSM should do if a modification occurs. This parameter is optional. The default value is SHRSTATIC.

Warning: Be careful about choosing a management class that contains a copy group that specifies SERIALIZATION=SHRDYNAMIC or SERIALIZATION=DYNAMIC.

If you choose a management class that allows a file to be archived while modifications are occurring, the archive copy stored on the server might be a fuzzy copy. A fuzzy copy is an archive copy that does not accurately reflect what is currently in the file. The fuzzy copy contains some, but not all, modifications. If that is not acceptable, be sure to choose a management class that allows ADSM to create an archive copy only if the file is not being modified.

If a file is restored that contains a fuzzy copy, the file may not be usable, depending on the file's application.

Possible values are:

SHRStatic

Specifies that a file must not be modified during archive. ADSM attempts to perform an archive as many as four times, depending on the value specified for CHANGINGRETRIES in the client user options file. If the file is modified during the archive attempt, ADSM does not archive the file.

Static

Specifies that a file must not be modified during archive. If a file is modified during the archive attempt, it is not archived.

Platforms that do not support the STATIC option default to SHRSTATIC. The STATIC option is not supported on the DOS platform.

SHRDynamic

Specifies that if the file is modified during an archive attempt, ADSM archives the file during its last attempt even though the file is being modified. ADSM attempts to perform an archive as many as four times, depending on the value specified for CHANGINGRETRIES in the client user options file.

Dynamic

Specifies that a file is archived on the first attempt by ADSM, regardless of whether the file is being modified during archive processing.

Examples

Task 1

Create a backup copy group named STANDARD for management class ACTIVEFILES in policy set VACATION in the EMPLOYEE_RECORDS policy domain. The backup destination is BACKPOOL, with a minimum interval of three days between backups, regardless of whether the files have been modified. Retain up to five backup versions while the file resides on the node's file system.

Command: DEFINE COPYGROUP employee_records vacation activefiles STANDARD Type=Backup DESTINATION=backpool FREQUENCY=3 VEREXISTS=5 MODE=ABSOLUTE

Output:

```
ANR1530I Backup copy group STANDARD defined in policy domain
EMPLOYEE_RECORDS, set VACATION, management class ACTIVEFILES.
```

Task 2

Create an archive copy group named STANDARD for management class MCLASS1 in policy set SUMMER in the PROG1 policy domain. The archive destination is ARCHPOOL, where the archive copy is kept for 730 days.

Command: DEFINE COPYGROUP prog1 summer mclass1 STANDARD Type=Archive DESTINATION=archpool RETVER=730

Output:

```
ANR1535I Archive copy group STANDARD defined in policy domain
PROG1, set SUMMER, management class MCLASS1.
```

Related Commands

Table 30 provides references to related commands.

Table 30. Commands Related to DEFINE COPYGROUP

Command	Description	Reference
ASSIGN DEFMGMTCLASS	Assigns an existing management class as the default for a specified policy set.	Page 40
COPY MGMTCLASS	Creates a copy of an existing management class.	Page 73
DEFINE MGMTCLASS	Defines a management class within the policy domain and policy set specified.	Page 114
DEFINE STGPOOL	Defines a storage pool as a named collection of data storage media.	Page 135
DELETE COPYGROUP	Deletes a backup or archive copy group from a policy domain and policy set.	Page 153
DELETE MGMTCLASS	Deletes a management class and its copy groups from a policy domain and policy set.	Page 172
EXPIRE INVENTORY	Manually starts inventory expiration processing.	Page 193
QUERY COPYGROUP	Displays the attributes of an existing copy group.	Page 270
QUERY MGMTCLASS	Displays information about one or more management classes.	Page 314
UPDATE COPYGROUP	Changes one or more attributes of a copy group.	Page 439

DEFINE DBBACKUPTRIGGER (Define the 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 using 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. When space usage reaches the percentage you specify, ADSM starts a background process that does the following:

- Runs a backup of the ADSM database using one or more scratch volumes.

Unless a full backup is required, ADSM runs an incremental backup. A full backup is required when any of the following conditions are true:

- The database has never been backed up.
- The database was loaded using a salvage dump created with the DSMSERV DUMPDB command and the database has not been backed up since it was loaded.
- The maximum number of incremental backups allowed between full backups has been reached.
- The log mode was changed from NORMAL to ROLLFORWARD.
- The database was extended or reduced.

ADSM can continue to write new records to the recovery log while creating a backup of the database.

Note: If a backup initiated by the database backup trigger fails, ADSM waits 60 seconds, and then tries again. It continues to retry the backup until it is successful. If the mount limit for the device class being used for the backup has been reached (all drives busy), ADSM automatically cancels lower priority operations (such as reclamation) to make a mount point available for the database backup. To ensure that a mount point is always available to run a database backup, you can define a device class that is used only by database backup.


- After the backup is successfully completed, ADSM deletes all recovery log records except those required to restore the database to its most current state when used with the backup of the database that was just created.

To display settings for the database backup trigger, use the QUERY DBBACKUPTRIGGER command. To update the settings for the database backup trigger, use the UPDATE DBBACKUPTRIGGER command. To delete the database backup trigger, use the DELETE DBBACKUPTRIGGER command.

Privilege Class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax



Parameters

DEVclass=*devclassname*

Specifies the name of the sequential access device class to use for backups. This parameter is required.

If the mount limit for the device class specified has already been reached (all drives are busy) when the backup runs, ADSM automatically cancels lower priority operations, such as reclamation, to make a mount point available for the backup.

Be sure that you have used the DEVCONFIG option in the server options file to specify an external file in which to store a backup copy of device class definitions. Otherwise, if your ADSM database is damaged or lost and must be restored, the definitions created using the DEFINE DEVCLASS command will not be available and will need to be recreated manually.

You should also use the VOLUMEHISTORY option in the server options file, to specify an external file in which to store a backup copy of sequential volume history information. The volume history file is also used for database recovery.

LOGFullpct=*percentage*

Specifies when ADSM automatically backs up the database, based on the percentage of available space used for the recovery log. When the percentage of space used for the recovery log exceeds this value, ADSM performs an automatic backup of the database and removes unnecessary records from the recovery log. This parameter is optional. You can specify a value from 1–99. The default value is 50 percent.

INCRDEVclass=*devclassname*

Specifies the name of the sequential access device class to use for incremental backups. This parameter is optional. If you do not specify a value for this option, all backups are run using the device class specified with the DEVCLASS parameter.

If the mount limit for the device class specified has already been reached (all drives are busy) when the backup runs, ADSM automatically cancels lower priority operations, such as reclamation, to make a mount point available for the backup.

Be sure that you have used the DEVCONFIG option in the server options file to specify an external file in which to store a backup copy of device class

definitions. Otherwise, if your ADSM database is damaged or lost and must be restored, the definitions created using the DEFINE DEVCLASS command will not be available and will need to be recreated manually.

NUMINCremental=number

Specifies the maximum number of incremental backups that can be run before a full backup is required. This parameter is optional. You can specify a value from 0–32. A 0 specifies that ADSM runs only full backups automatically. The default value is 6.

Each incremental backup, whether run automatically by ADSM or run by issuing the BACKUP DB command, is added to the count of incremental backups run.

Each full backup, whether run automatically by ADSM or run by issuing the BACKUP DB command, resets the count for incremental backups to zero.

Examples

Task

Define the setting of the database backup trigger at 80 percent. Run two incremental backups to every full backup of the database. Assume a device class of FILE for the backup.

Command: DEFINE DBBackuptrigger DEVclass=file LOGFullpct=80
NUMINCremental=2

Output:

```
ANR2283I Database backup trigger defined, but is disabled.
```

Related Commands

Table 31 provides references to related commands.

Table 31. Commands Related to DEFINE DBBACKUPTRIGGER

Command	Description	Reference
BACKUP DB	Backs up the ADSM database to sequential access volumes.	Page 50
DEFINE DEVCLASS	Defines a device class that can be used for storage pools, importing, or exporting.	Page 99
DELETE DBBACKUPTRIGGER	Deletes the database backup trigger.	Page 156
QUERY DBBACKUPTRIGGER	Displays information about the database backup trigger.	Page 278
SET LOGMODE	Sets the mode for saving records in the recovery log.	Page 416
UPDATE DBBACKUPTRIGGER	Changes one or more attributes of the database backup trigger.	Page 447

DEFINE DBCOPY (Define a Volume Copy of a Database Volume)

Use the DEFINE DBCOPY command to create a volume copy of an existing database volume. Database volume copies increase the availability of the database and provide safeguard against media failure. You can have as many as two volume copies of an existing database volume. After each database volume copy is defined, the server attempts to synchronize the volume copy with the database.

All database volumes must reside on random access volumes. Before issuing this command, you must allocate and format the volume by using the ANRFMT job or the ANRFMT2 job found in ADSM.SAMPLIB. For information on using these jobs, see *ADSM Administrator's Guide*.

Volume copies must have at least the same capacity as the original volume and should be defined on separate physical devices. Any additional space on the volume copy is not used.

This command can generate a background process that can be canceled with the CANCEL PROCESS command. If a DEFINE DBCOPY background process is canceled, the volume copy is created; however, synchronization processing is canceled, leaving the volume in a *stale* state. To display information on background processes, use the QUERY PROCESS command.

Privilege Class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
►►—DEFine DBCopy—volname—copyvolname—————►►
```

Parameters

volname

Specifies the name of the database volume for which you want to create a volume copy. This parameter is required.

copyvolname

Specifies the name of the volume you want to assign as a database volume copy. This parameter is required.

For VSAM linear data sets, specify a standard VSAM cluster name. A VSAM cluster name consists of one or more qualifiers, which are delimited by periods. The length of each qualifier is from 1 to 8 alphanumeric characters. The maximum length of the cluster name is 42 characters. An example is ADSM.BFS.POOL2.VOLA.

Any volume name used by ADSM cannot contain any embedded blanks or equal signs.

Examples

Task

Create a volume copy of the database volume named HARVEY.ADSM.PINV with the volume named GLEDDIE.ADSM.PINV.

Command: DEFine DBCopy harvey.adsm.pinv gleddie.adsm.pinv

Output:

```
ANR2242I Database volume copy GLEDDIE.ADSM.PINV defined;
synchronization process started (process id 1).
ANS5104I Process number 1 started.
```

Related Commands

Table 32 provides references to related commands.

Table 32. Commands Related to DEFINE DBCOPY

Command	Description	Reference
CANCEL PROCESS	Cancels a background server process.	Page 64
DEFINE DBVOLUME	Assigns a volume to be used for database storage.	Page 97
DELETE DBVOLUME	Deletes a volume from the database.	Page 158
QUERY DBVOLUME	Displays information about one or more database volumes.	Page 280
QUERY PROCESS	Displays information about one or all background server processes.	Page 332

DEFINE DBVOLUME (Define a Database Volume)

Use the DEFINE DBVOLUME command to define a new database volume. All database volumes must reside on random access volumes.

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

Before issuing this command, you must allocate and format the volume using the ANRFMT job or the ANRFMT2 job found in ADSM.SAMPLIB. For information on using these jobs, see *ADSM Administrator's Guide*.

When you define a database volume, it must occupy at least 5MB of space. Regardless of the defined storage capacity of a database volume, 1MB is not available for database storage. For example, if you define a database volume with a capacity of 9MB, 8MB are available for storage on the volume.

Each time you define a new database volume, you provide additional space to the database. To permit the database to use this additional space, use the EXTEND DB command.

Privilege Class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
►►—DEFine DBVolume—volname—————◄◄
```

Parameters

volname

Specifies the name of the database volume to be defined. This parameter is required. For VSAM linear data sets, specify a standard VSAM cluster name. A VSAM cluster name consists of one or more qualifiers that are separated by periods. The length of each qualifier is from 1 to 8 alphanumeric characters. The maximum length of the cluster name is 42 characters. An example is ADSM.BFS.POOL2.VOLA.

Volume names used by ADSM cannot contain any embedded blanks or equal signs.

DEFINE DBVOLUME

Examples

Task

Create a database volume named COY.ADSM.DBVOL

Command: DEFine DBVolume coy.adsm.dbvol

Output:

```
ANR2240I Database volume COY.ADSM.DBVOL defined.
```

Related Commands

Table 33 provides references to related commands.

Table 33. Commands Related to DEFINE DBVOLUME

Command	Description	Reference
DEFINE DBCOPY	Creates a volume copy of a database volume.	Page 95
DELETE DBVOLUME	Deletes a volume from the database.	Page 158
EXTEND DB	Extends usable database storage within the volumes previously allocated for its use.	Page 213
QUERY DBVOLUME	Displays information about one or more database volumes.	Page 280

DEFINE DEVCLASS (Define a Device Class)

Use the DEFINE DEVCLASS command to define a device class for sequential access devices that can be used for storage pools.

You can use the UPDATE DEVCLASS and DELETE DEVCLASS commands to either modify or eliminate existing device classes. However, you cannot modify the device type of an existing device class. The DISK device class is predefined by ADSM and requires no user intervention. This also means that the DISK device class cannot be modified or eliminated via the UPDATE DEVCLASS or DELETE DEVCLASS commands. The CARTRIDGE device class is predefined by ADSM and can be modified.

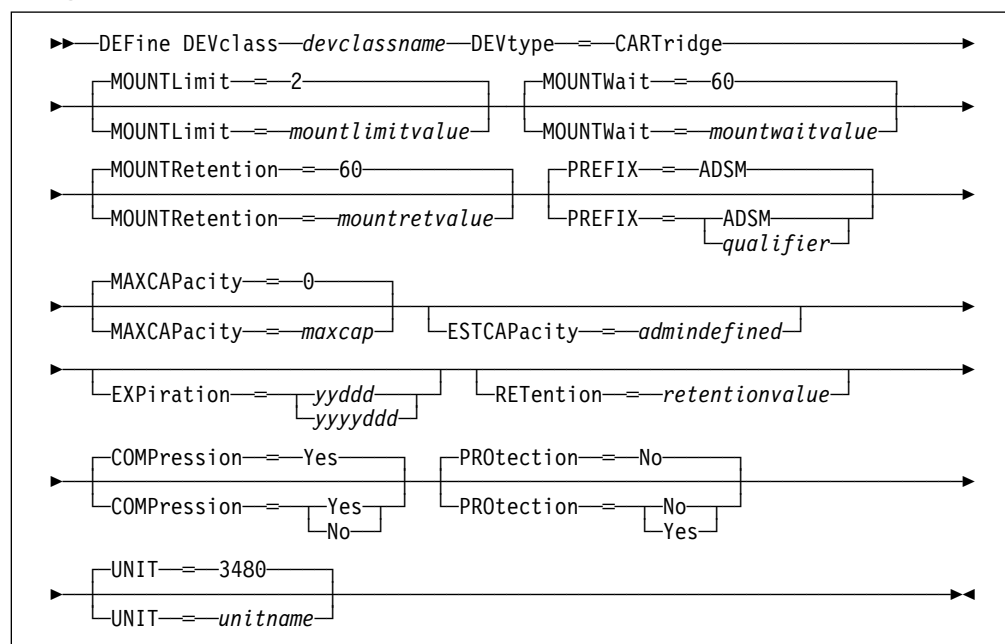
You should define a unique device class to support the different tape devices, tape recording technologies, and storage media available at your installation. For more information on defining device classes, see the *ADSM Administrator's Guide*.

If you include the DEVCONFIG option in the server options file, the files you specify with that option are automatically updated with the results of this command. When you use this option, the files specified are automatically updated whenever a device class is defined, updated, or deleted.

Privilege Class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax for Defining a Device Class with a Device Type of CARTRIDGE



Parameters for Defining a Device Class With a Device Type of CARTRIDGE

devclassname

Specifies the name of the device class to be defined. This parameter is required. The maximum length of the device class name is 30 characters.

DEVtype=CARTRIDGE

Specifies the CARTRIDGE device type is assigned to the device class. This parameter is required. CARTRIDGE indicates that cartridge tape devices are assigned to the device class.

MOUNTLimit=mountlimitvalue

Specifies the maximum number of tape volumes that can be simultaneously mounted for a device class. This parameter is optional. The default value is 2. You can specify an integer in the range of 1 to 256, inclusive.

If this value is set to 1, then it may not be possible for tape-to-tape operations (like reclamation) to work.

You cannot share drives between device classes (that is, the sum of your mount limits should not be more than the number of physical drives allocated to ADSM).

MOUNTWait=mountwaitvalue

Specifies the maximum length of time, in minutes, to wait for a tape volume to be mounted. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. The default value is 60 minutes. You can specify an integer in the range of 0 to 9999, inclusive.

MOUNTRetention=mountretvalue

Specifies the amount of time, in minutes, to retain an idle tape volume before dismounting it. Mountretention time begins after the idle time-out period has expired. This parameter is optional. The default value is 60. You can specify an integer in the range of 0 to 9999, inclusive.

This parameter can improve response time for tape mounts by leaving the previously mounted volume online.

PREFIX=tapevolprefix

Specifies the high-level qualifier of the data set name in the tape labels. For all tape volumes assigned to this device class, ADSM uses this prefix to create the data set name. This parameter is optional. The default value is ADSM.

An example of a tape volume data set name using the default prefix is ADSM.BFS. You can specify any valid MVS data set qualifier for the prefix.

If you have already established a tape label naming convention that supports your current tape management system, consider using a tape volume prefix that conforms to your naming conventions.

MAXCAPacity=maxcap

Specifies the maximum amount of data ADSM can store in the volumes categorized by this device class. This parameter is optional. The default is 0.

The value you specify for this parameter (other than 0) is used to determine maximum capacity of the volume. If you do not specify a value and 0 is used, ADSM uses a default method based on IBM 3480 and 3490E device charac-

teristics using IBM tapes. If this default method does not meet your needs, set the value to a nonzero value less than the maximum capacity of the tape to avoid an end-of-tape situation. If ADSM detects an end-of-tape situation for the specified volume, a warning message is issued and the tape is treated as full.

You must specify this value as an integer followed by a scale factor. The following are possible scale factors:

Symbol	Symbol Name
K	Kilobyte
M	Megabyte
G	Gigabyte

For example, MAXCAPACITY=5M specifies that the maximum capacity for this device class is 5 megabytes.

ESTCAPacity=admindefined

Specifies the estimated capacity for volumes categorized by this device class. This parameter is optional.

This value is ignored if you specify a value other than 0 in the MAXCAPACITY parameter. If ADSM uses the value specified in the ESTCAPACITY parameter, this value does not determine the amount of data stored on the volume, but only determines the usage before the volume is filled. After a volume is full, the actual amount of data stored on the tape is used for these calculations.

You must specify this value as an integer followed by a scale factor. Possible scale factors are:

Symbol	Symbol Name
K	Kilobyte
M	Megabyte
G	Gigabyte

For example, ESTCAPACITY=5M specifies that the estimated capacity for a volume in this device class is 5 megabytes. The smallest value allowed is 100KB (that is, ESTCAPACITY=100K).

For more information on the default estimated capacity for cartridge tapes, see Table 34.

Table 34. ADSM Estimated Capacity for Cartridge Tape

Tape Format	Media Type	Estimated Capacity
3480	(Not applicable)	180MB
3480XF	(Not applicable)	180MB
3490-A01	CST	180MB
3490B	CST	360MB
3490C	CST	360MB
3490B	ECCST	720MB
3490C	ECCST	720MB

Note:

3490 Base devices include models A01, A02, A10, A20, B02, B04, B20, B40, D41, D42
 3490 E devices include models C1A, C2A, C10, C11, C22

EXPIRATION=*expirationdate*

Specifies the expiration date that is placed on the tape labels for this device class, which are used for allocation. This parameter is optional. There is no default value.

If you specify this parameter, enter the expiration date when ADSM no longer requires the tape. ADSM does not use this information, however; this information is passed to the operating system for use by MVS or tape management systems. For additional information, see the *ADSM Administrator's Guide*.

If you specify this parameter, you cannot specify the RETENTION parameter. Possible values are:

yyddd

Specifies the expiration date for this device class, in YYDDD format.

yyyyddd

Specifies the expiration date for this device class, in YYYYDDD format.

RETENTION=*retentionvalue*

Specifies an integer that defines the length of time, in days, to retain the tape. This parameter is optional. You can specify an integer in the range of 1 to 9999, inclusive. If you specify this parameter, enter the number of days that equals the length of time ADSM is expected to use the tape. ADSM does not use this information, however; this information is passed to the operating system for use by MVS or tape management systems. For additional information, see the *ADSM Administrator's Guide*.

If you specify this parameter, you cannot specify the EXPIRATION parameter.

COMPRESSION=*compressionvalue*

Specifies whether the device class is to exploit file compression. This parameter is optional. The default value is YES.

If your tape devices have IDRC, then data can be compressed when written to CARTRIDGE tape devices. If your tape devices do not have IDRC, then the compression attribute is ignored and data is not compressed.

Also, if your tape volumes can be moved to other physical sites, ensure that the compressed data can be read at the other sites. If the other sites cannot read compressed data, you want to specify COMPRESSION=NO. Possible values are:

Yes

Specifies that the data for each tape volume is compressed.

No

Specifies that the data for each tape volume is not compressed.

PROTECTION=*protectionvalue*

Specifies whether RACF, if installed, protects all volumes assigned to this device class. This parameter is optional. The default value is NO. Possible values are:

No

Specifies that RACF does not protect all CARTRIDGE volumes assigned to this tape device.


Yes

Specifies that RACF protects all CARTRIDGE volumes assigned to this tape device.

UNIT=unitname

Specifies an esoteric unit name to specify a group of tape devices that support Cartridge System tape. This parameter is optional. The default value is 3480. This name can be as many as 8 characters.

Syntax for Defining a Device Class with a Device Type of FILE



Parameters for Defining a Device Class With a Device Type of FILE

devclassname

Specifies the name of the device class to be defined. This parameter is required. The maximum length of the device class name is 30 characters.

DEVTYPE=FILE

Specifies the FILE device type is assigned to the device class. This parameter is required.

FILE indicates that a file or files are assigned to this device class. When the server needs to access a volume that belongs to this device class, it opens a file and reads or writes file data.

MOUNTLIMIT=mountlimitvalue

Specifies the maximum number of FILE volumes that can be concurrently open at any one time for this device class. This parameter is optional. The default value is 2. For 3995 devices emulating 3390 devices, this value should not be set higher than the number of concurrent input/output streams possible on the media storing the ADSM volumes.

Because FILE volumes are classified as sequential media, processing conventions normally associated with sequential media are followed. Specifically, this means that mount limits are observed, and that data that has expired leaves *holes* in the volume and volumes can only be recycled through either the MOVE DATA command or the volume reclamation process. Only one process can read or write to a volume at a time. This has important implications for clients backing up directly to FILE media. The number of active clients is limited to the value specified in the MOUNTLIMIT parameter for the device class.

The value you specify in this parameter is important if there is a significant penalty switching from one volume to another (switching can take place when using 3995 devices emulating 3390 devices). The value you specify in this parameter should be no higher than the number of physical drives available on the device.

See the Usage Notes for more information on FILE volumes.

DEFINE DEVCLASS

PREFIX=*qualifier*

Specifies the MVS high-level qualifier that is used to allocate scratch volume data sets. It must be 1–8 alphanumeric characters and begin with an alpha. The default is ADSM.

MAXCAPacity=*maxcap*

Specifies the maximum capacity of the FILE volumes. This parameter is optional. The default is 1MB. The largest accepted value is 2047MB (2GB minus 1MB). The smallest accepted value is 100KB.

Volume capacity is determined by the value you specify for the device class. If you are explicitly allocating FILE volume data sets, be sure the allocated capacity matches the MAXCAPACITY parameter. FILE volumes are allocated in 16KB blocks. Thus, if the MAXCAPACITY value is 50MB, you should allocate 3200 blocks (16KB per block). A combination of primary and secondary extents may be specified to achieve the allocation.

Note: Note that FILE volumes will only be filled to the amount specified by the MAXCAPACITY parameter. Thus, if a larger data set is explicitly allocated, the additional space will not be used.

If the data set is allocated to a size smaller than the value you specify for the MAXCAPACITY parameter, writing will stop after all allocated blocks have been used. If this situation occurs, a message is issued, the volume status remains "Filling," and access is set to read-only.

You must specify this value as an integer followed by a scale factor. Possible scale factors are:

Symbol	Symbol Name
K	Kilobyte
M	Megabyte
G	Gigabyte

For example, MAXCAPACITY=5M specifies that the maximum capacity for this device class is 5 megabytes.

Usage Notes

- The 3490-A01 is treated like a 3480. The capacity shown in a QUERY VOLUME will, therefore, reflect 3480 capacities.
- The estimated capacity of a tape volume is determined by the format, length, and density of the volume.
- The format for a tape volume is determined by the device type and the recording technology of the tape volume.
- The EXPIRATION and RETENTION parameters are mutually exclusive; if you specify one, you cannot specify the other.
- FILE device class volumes are represented by sequential MVS data sets. There are two primary methods for allocating FILE volumes: scratch and explicit allocation.

ADSM allocates scratch volumes when you specify the MAXSCRATCH parameter on a storage pool allocated from a FILE device class. ADSM generates the name of your data set using a combination of the value you specified in the PREFIX parameter as the HLQ (high-level qualifier) and a system-generated extension.

Judicious use of the PREFIX parameter can assist the installation in identifying and isolating ADSM FILE volumes.

You can also use the DEFINE VOLUME command to explicitly define volumes to storage pools from a FILE device class. The volume name you specify is used as the fully qualified data set representing the volume. Explicitly defined volumes are used before scratch volumes. The data set representing the explicitly defined volume may or may not be allocated ahead of time. If it is not allocated in advance, the first time write processing occurs ADSM attempts to dynamically allocate the file.

The QUERY VOLUME command lists the fully qualified data set name corresponding to the FILE volume.

- ADSM attempts to dynamically allocate a data set for any volume not explicitly allocated ahead of time (this would be the usual way of allocating scratch volumes). If the capacity of a volume is 256MB or less, a single primary extent (0 secondary extents) is allocated. For volumes greater than 256MB, a primary extent of 256MB is allocated, with secondary extents of 1/15th of the remaining required space.

Explicitly defined volumes may either be defined before they are used or dynamically allocated. If the volume is predefined, it should have the following characteristics:

```

Data Set Name: ADSM.FILEVOL0

          General Data:
Volume serial:      volser
Device type:        3390
Organization:       PS
Record format:      F
Record length:      16384
Block size:         16384
* First extent blocks: 16386
* Secondary blocks:  1093

Creation date:      1995/02/16
Expiration date:    ***NONE***
    
```

* These values vary with volume size.

- FILE volume reclamation uses the same methodology as other sequential media. The key factor for recycling a reclaimed volume is whether the volume was automatically defined by ADSM as a scratch volume or explicitly defined by an ADSM administrator using the DEFINE VOLUME command. Explicitly defined volumes are marked as empty and available for subsequent mounts at the completion of reclamation. Scratch volumes are deleted from the storage pool in which it was defined and the MVS data set corresponding to the scratch volume is automatically erased from the system.

Examples

Task

Define a device class named SQUARE1 with a CARTRIDGE tape device and with RACF protection active for all tape volumes assigned to this device class. All data is compressed for this device class.

Command: DEFine DEVclass square1 DEVtype=CARTridge COMPression=Yes
PROtection=Yes

Output:

```
ANR2203I Device class SQUARE1 defined.
```

Related Commands

Table 35 provides references to related commands.

Table 35. Command Related to DEFINE DEVCLASS

Command	Description	Reference
BACKUP DEVCONFIG	Backs up ADSM device information to a file.	Page 54
DEFINE VOLUME	Assigns a volume to be used for storage within a specified storage pool.	Page 148
DELETE DEVCLASS	Deletes a device class name.	Page 160
QUERY DEVCLASS	Displays information about one or more device classes.	Page 283
QUERY VOLUME	Displays information about one or more storage pool volumes.	Page 359
UPDATE DEVCLASS	Changes the attributes of a device class.	Page 450

DEFINE DOMAIN (Define a New Policy Domain)

Use the DEFINE DOMAIN command to define a new policy domain.

A policy domain provides you with a logical way of managing backup and archive policies for a group of client nodes with common backup-archive policy needs. You can also define policy domains to limit the number of client nodes that can be managed by a single policy administrator. A policy domain contains one or more policy sets, with only one policy set active at a time.

Each policy set contains one or more management classes. You can specify one default management class for each policy set. A management class can contain a backup copy group, an archive copy group, or both; however, neither is required.

An ADSM client node is assigned to one policy domain. The active policy set within the policy domain should contain management classes that meet the needs of files stored on the client node. The user of the client node can use the default management class or select any management class in the active policy set to manage files that are backed up or archived from the client node.

This command specifies the backup and archive retention grace periods for the policy domain. These retention periods are used by ADSM to protect backup versions and archive copies from deletion when the respective copy group is not available. For further descriptions of these parameters, see the parameter definitions. You can also find a detailed discussion of the retention grace periods in the *ADSM Administrator's Guide*.

The backup and archive retention values are used by ADSM to manage inactive copies of files when the management class to which a file is bound no longer exists, and the default management class does not contain a backup copy group.

To determine the retention periods for an existing policy domain, use the QUERY DOMAIN command.

Client nodes running on different platforms can be assigned to the same policy domain. For example, a policy domain can be used by both OS/2 and AIX clients.


After you define a new policy domain, you can define additional policy objects in the new policy domain, such as policy sets, management classes, and copy groups. You can use COPY commands to copy policy sets and management classes. In addition, you must activate a policy set before clients assigned to the new policy domain can back up or archive files. A new policy domain can also be created by copying an existing policy domain to a new policy domain name.

Privilege Class

To issue this command, you must have system privilege.

DEFINE DOMAIN

Syntax



Parameters

domainname

Specifies the name of the policy domain to be defined. This parameter is required. The maximum length of this name is 30 characters.

DESCRIPTION=*description*

Specifies a text string that describes the policy domain. This parameter is optional. The maximum length of the description is 255 characters. The description must be enclosed in quotation marks if it contains any blank characters. The default is no description.

BACKRETENTION=*bkretvalue*

Specifies an integer that defines the number of days (from the date of deactivation) to retain backup versions that are no longer on the client's system. This parameter is optional. The default value is 30. You can specify an integer in the range of 0 to 9999, inclusive. The backup retention value is used by ADSM to manage inactive copies of files when either of the following conditions are present:

- A file's management class is changed, but neither the new management class nor the default management class contains a backup copy group.
- The management class to which a file is bound no longer exists, and the default management class does not contain a backup copy group.

ARCHRETENTION=*archretvalue*

Specifies an integer that defines the number of days (from the date of archive) to retain archive copies. This parameter is optional. The default value is 365. You can specify an integer in the range of 0 to 9999, inclusive. The archive retention value is used by ADSM to manage inactive copies of files when either of the following conditions are present:

- A file's management class is changed, but neither the new management class or default management class contains an archive copy group.
- The management class to which a file is bound no longer exists, and the default management class does not contain an archive copy group.

Examples

Task 1

Define a policy domain with the name of EMPLOYEE_RECORDS.

Command: DEFine D0main employee_records

Output:

```
ANR1500I Policy domain EMPLOYEE_RECORDS defined.
```

Task 2

Define a policy domain with a name of PROG1 and a description of Programming Group Domain. Specify that archive copy retention is for 90 days when management classes or archive copy groups are deleted and the default management class does not contain an archive copy group. Also specify that backup version retention is for 60 days when management classes or copy groups are deleted and the default management class does not contain a backup copy group.

Command: DEFine D0main prog1 DESCription="Programming Group Domain"
BACKRETention=60 ARCHRETention=90

Output:

```
ANR1500I Policy domain PROG1 defined.
```

Related Commands

Table 36 provides references to related commands.

Table 36. Commands Related to DEFINE DOMAIN

Command	Description	Reference
COPY DOMAIN	Creates a copy of an existing policy domain.	Page 70
DEFINE POLICYSET	Defines a policy set within the specified policy domain.	Page 118
DELETE DOMAIN	Deletes a policy domain along with any policy objects in the policy domain.	Page 162
QUERY DOMAIN	Displays information about one or more policy domains.	Page 287
UPDATE DOMAIN	Changes the attributes of a policy domain.	Page 458

DEFINE LOGCOPY (Define a Volume Copy of a Recovery Log Volume)

Use the DEFINE LOGCOPY command to create a volume copy of an existing recovery log volume. Recovery log volume copies increase the availability of the recovery log and provide a safeguard against media failure. You can have as many as two volume copies of an existing recovery log volume. After each recovery log volume copy is defined, the server attempts to synchronize the volume copy with the recovery log.

All recovery log volumes must reside on random access volumes. Before issuing this command, you must allocate and format the volume by using the ANRFMT job or the ANRFMT2 job found in ADSM.SAMPLIB. For information on using these jobs, see *ADSM Administrator's Guide*.

Volume copies must have at least the same capacity as the original volume and should be defined on separate physical devices. Any additional space on the volume copy is not used.

This command can generate a background process that can be canceled with the CANCEL PROCESS command. If a DEFINE LOGCOPY background process is canceled, the volume copy is created; however, synchronization processing is canceled leaving the volume in a *stale* state. To display information on background processes, use the QUERY PROCESS command.

Privilege Class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
►►—DEFine LOGCopy—volname—copyvolname—◄◄
```

Parameters

volname

Specifies the name of the recovery log volume for which you want to create a volume copy. This parameter is required.

copyvolname

Specifies the name of the volume you want to assign as a recovery log volume copy. This parameter is required. For VSAM linear data sets, specify a standard VSAM cluster name. A VSAM cluster name consists of one or more qualifiers, which are delimited by periods. The length of each qualifier is from 1 to 8 alphanumeric characters. The maximum length of the cluster name is 42 characters. An example is ADSM.BFS.SPOOL2.VOLA.

Volume names used by ADSM cannot contain any embedded blanks or equal signs.

Examples

Task

Create a volume copy of the recovery log volume named HARVEY.ADSM.PLOG with the volume named GLEDDIE.ADSM.PLOG.

Command: DEFine LOGCopy harvey.adsm.plog gleddie.adsm.plog

Output:

```
ANR2262I Recovery log volume copy GLEDDIE.ADSM.PLOG
defined; synchronization process started (process id 2).
ANS5104I Process number 2 started.
```

Related Commands

Table 37 provides references to related commands.

Table 37. Commands Related to DEFINE LOGCOPY

Command	Description	Reference
CANCEL PROCESS	Cancels a background server process.	Page 64
DEFINE LOGVOLUME	Defines a disk volume to the recovery log.	Page 112
DELETE LOGVOLUME	Deletes a volume from the recovery log.	Page 170
QUERY LOGVOLUME	Displays information about one or more log volumes.	Page 311
QUERY PROCESS	Displays information about one or all background server processes.	Page 332

DEFINE LOGVOLUME (Define a Recovery Log Volume)

Use the DEFINE LOGVOLUME command to define a new recovery log volume. All recovery log volumes must reside on random access volumes.

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

Before issuing this command, you must allocate and format the volume by using the ANRFMT job or the ANRFMT2 job found in ADSM.SAMPLIB. For information on using these jobs, see *ADSM Administrator's Guide*.

When you define a recovery log volume, it must occupy at least 5MB of space. Regardless of the defined storage capacity of a recovery log volume, 1MB is not available for recovery log storage. For example, if you define a recovery log volume with a capacity of 9MB, 8MB are available for storage on the volume.

Each time you define a new recovery log volume, you provide additional space for the recovery log. To permit the recovery log to use this additional space, use the EXTEND LOG command.

Privilege Class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
►►—DEFine LOGVolume—volname—————◄◄
```

Parameters

volname

Specifies the name of the recovery log volume. This parameter is required. For VSAM linear data sets, specify a standard VSAM cluster name. A VSAM cluster name consists of one or more qualifiers, which are delimited by periods. The length of each qualifier is from 1 to 8 alphanumeric characters. The maximum length of the cluster name is 42 characters. An example is ADSM.BFS.POOL2.VOLA.

Any volume name used by ADSM cannot contain any embedded blanks or equal signs.

Examples

Task

Define a recovery log volume named GLEDDIE.ADSM.PLOG.

Command: DEFine LOGVolume gleddie.adsm.plog

Output:

```
ANR2260I Recovery log volume GLEDDIE.ADSM.PLOG defined.
```

Related Commands

Table 38 provides references to related commands.

Table 38. Commands Related to DEFINE LOGVOLUME

Command	Description	Reference
DEFINE LOGCOPY	Creates a volume copy of a recovery log volume.	Page 110
DELETE LOGVOLUME	Deletes a volume from the recovery log.	Page 170
EXTEND LOG	Extends usable log storage within the volumes previously allocated for its use.	Page 215
QUERY LOGVOLUME	Displays information about one or more log volumes.	Page 311

DEFINE MGMTCLASS (Define a Management Class)

Use the DEFINE MGMTCLASS command to define a new management class in a policy set.

You can define one or more management classes for each policy set in a policy domain. A management class can contain a backup copy group, an archive copy group, or both. However, neither is required.

Note: The DEFINE MGMTCLASS command will fail if a copy storage pool is specified as the destination for space-managed files.

You must specify an existing management class as the default management class for a policy set before you can activate the policy set. To specify an existing management class as the default management class, use the ASSIGN DEFMGMTCLASS command.

When a client node is registered with ADSM, the node is defined to a policy domain. This policy domain contains management classes appropriate for the information stored on the client node.

The user of a client node can use the default management class or select any management class in the active policy set. By entering clear descriptions of your management classes, users can easily select management classes that are appropriate for their storage management needs.

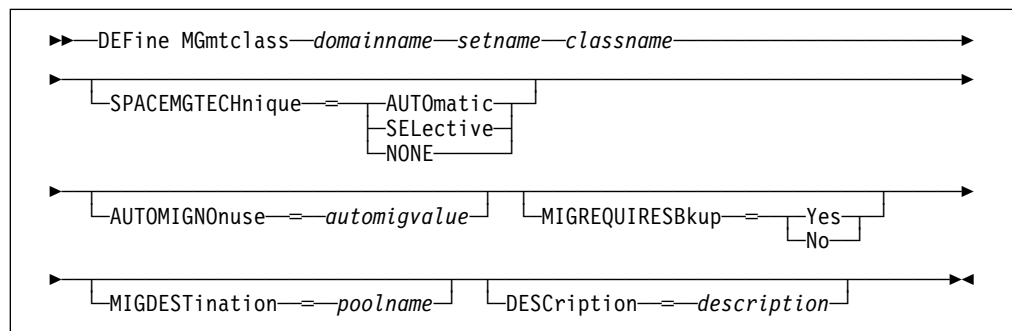
The following parameters affect only HSM clients that use this management class, and have no effect on backup-archive clients:

- SPACEMGTECHNIQUE
- AUTOMIGNONUSE
- MIGREQUIRESBKUP
- MIGDESTINATION

Privilege Class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the management class belongs.

Syntax



Parameters

domainname

Specifies the name of the policy domain to which the management class belongs. This parameter is required.

setname

Specifies the name of the policy set to which the management class belongs. This parameter is required. You cannot define a management class to the ACTIVE policy set.

classname

Specifies the name of the new management class. This parameter is required. The maximum length of this name is 30 characters. You cannot use either *default* or *grace_period* as a class name.

SPACEMGTECHnique=*smtvalue*

Specifies whether the file is eligible for migration. This parameter is optional. Possible values are:

AUTOMATIC

Specifies that the file is eligible for automatic migration or selective migration.

SElective

Specifies that the file is eligible for selective migration only.

NONE

Specifies that the file is not eligible for migration.

The default in the management class named STANDARD is NONE.

This parameter is effective only for HSM clients, not for backup-archive clients.

AUTOMIGNOnuse=*automigvalue*

Specifies the number of days that must elapse since the file was last accessed before it is eligible for automatic migration. This parameter is optional. If SPACEMGTECHNIQUE=AUTOMATIC is not specified, ADSM ignores this attribute. You can specify a value from 0–9999. The default value is 0.

This parameter is effective only for HSM clients, not for backup-archive clients.

MIGREQUIRESBkup=*migreqvalue*

Specifies whether a backup version of a file must exist before a file can be migrated. This parameter is optional. Possible values are:

Yes A backup version must exist.

No A backup version is optional.

This parameter is effective only for HSM clients, not for backup-archive clients.

MIGDESTination=*poolname*

Specifies the name of a primary storage pool that is a destination for migrated files. A file must be associated with a management class that specifies an appropriate storage pool for migrated files. The type of storage devices in a

DEFINE MGMTCLASS

storage pool and options specified for the storage pool affects the speed with which a migrated file can be recalled to a local file system. This parameter is optional. The default value is SPACEMGPOOL.

This parameter is effective only for HSM clients, not for backup-archive clients.

DEScRiption=*description*

Describes the management class by using a text string. This parameter is optional. The maximum length of the description is 255 characters. The description must be enclosed in quotation marks if it contains any blank characters.

Examples

Task 1

Define a management class called ACTIVEFILES for policy set VACATION in the EMPLOYEE_RECORDS policy domain.

Command: `DEFine Mgmtclass employee_records vacation activefiles`

Output:

```
ANR1520I Management class ACTIVEFILES defined in policy domain EMPLOYEE_RECORDS, set VACATION.
```

Task 2

Define a management class called MCLASS1, described as Technical Support Mgmt Class, for policy set SUMMER in the PROG1 policy domain.

Command: `DEFine Mgmtclass prog1 summer mclass1 DEScRiption="Technical Support Mgmt Class"`

Output:

```
ANR1520I Management class MCLASS1 defined in policy domain PROG1, set SUMMER.
```

Related Commands

Table 39 provides references to related commands.

Table 39 (Page 1 of 2). Commands Related to DEFINE MGMTCLASS

Command	Description	Reference
ASSIGN DEFMGMTCLASS	Assigns an existing management class as the default for a specified policy set.	Page 40
COPY MGMTCLASS	Creates a copy of an existing management class.	Page 73

Table 39 (Page 2 of 2). Commands Related to DEFINE MGMTCLASS

Command	Description	Reference
DEFINE COPYGROUP	Defines a copy group for backup or archive processing within a specified management class.	Page 83
DEFINE POLICYSET	Defines a policy set within the specified policy domain.	Page 118
DELETE MGMTCLASS	Deletes a management class and its copy groups from a policy domain and policy set.	Page 172
QUERY COPYGROUP	Displays the attributes of an existing copy group.	Page 270
QUERY MGMTCLASS	Displays information about one or more management classes.	Page 314
QUERY POLICYSET	Displays information about one or more policy sets.	Page 329
UPDATE COPYGROUP	Changes one or more attributes of a copy group.	Page 439
UPDATE MGMTCLASS	Changes the attributes of a management class.	Page 460

DEFINE POLICYSET (Define a Policy Set)

Use the DEFINE POLICYSET command to define a policy set in a specified policy domain.

A policy set provides you with a logical way of grouping management classes and a set of definitions for use in the ACTIVE policy set. You can define one or more policy sets for each policy domain; however, only one policy set can be active at a time in a given policy domain.

A policy set can contain zero or more management classes. You can specify one default management class for each policy set. A management class can contain a backup copy group, an archive copy group, or both. However, neither is required.


Use the ACTIVATE POLICYSET command to make a new policy set the active policy set. When a policy set is activated, ADSM uses these new definitions to manage client files for the client nodes assigned to the policy domain where this newly activated policy set resides. Use the VALIDATE POLICYSET command to verify that a policy set is complete and valid before activating it with the ACTIVATE POLICYSET command.

After you define a policy set, you can define policy objects to the new policy set, such as management classes, a default management class, and copy groups. You can use the COPY MGMTCLASS command to copy a management class and its associated copy groups into the new policy set.

Privilege Class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the policy set belongs.

Syntax



Parameters

domainname

Specifies the name of the policy domain to which the policy set belongs. This parameter is required.

setname

Specifies the name you want to assign to the policy set. This parameter is required. The maximum length of this name is 30 characters. You cannot define a policy set named ACTIVE.

DESCRIPTION=description

Describes the new policy set using a text string. This parameter is optional. The maximum length of the description is 255 characters. The description must be enclosed in quotation marks if it contains any blank characters. The default is no description.

Examples

Task 1

Define a policy set called VACATION for the EMPLOYEE_RECORDS policy domain.

Command: DEFine POLicyset employee_records vacation

Output:

```
ANR1510I Policy set VACATION defined in policy domain
EMPLOYEE_RECORDS.
```

Task 2

Define a policy set called SUMMER for the PROG1 policy domain and include a description of Programming Group Policies.

Command: DEFine POLicyset prog1 summer DESCription="Programming Group Policies"

Output:

```
ANR1510I Policy set SUMMER defined in policy domain PROG1.
```

Related Commands

Table 40 provides references to related commands.

Table 40 (Page 1 of 2). Commands Related to DEFINE POLICYSET

Command	Description	Reference
ACTIVATE POLICYSET	Validates and activates a policy set within a policy domain.	Page 38
COPY MGMTCLASS	Creates a copy of an existing management class.	Page 73
COPY POLICYSET	Creates a copy of an existing policy set.	Page 75
DEFINE DOMAIN	Defines a policy domain that can be assigned to one or more clients.	Page 107
DEFINE MGMTCLASS	Defines a management class within the policy domain and policy set specified.	Page 114
DELETE POLICYSET	Deletes a policy set, including its management classes and copy groups, from a policy domain.	Page 174

DEFINE POLICYSET

Table 40 (Page 2 of 2). Commands Related to DEFINE POLICYSET

Command	Description	Reference
QUERY POLICYSET	Displays information about one or more policy sets.	Page 329
UPDATE POLICYSET	Changes the description of a policy set.	Page 467
VALIDATE POLICYSET	Verifies and reports on conditions the administrator must consider before activating the policy set.	Page 501

DEFINE SCHEDULE (Define a Client or an Administrative Command Schedule)

Use the DEFINE SCHEDULE command to create a new schedule. The DEFINE SCHEDULE command takes two forms, depending on whether the schedule applies to client operations or administrative commands. The syntax and parameters for each form are defined separately.

Defining Client Schedules

Use the DEFINE SCHEDULE command to create a new schedule for client operations such as backing up or archiving clients in a specified policy domain.

A schedule determines:

- The automatic function to be performed, for example, an incremental backup
- A startup window that indicates a time range during which the operation can begin
- The time period between startup windows

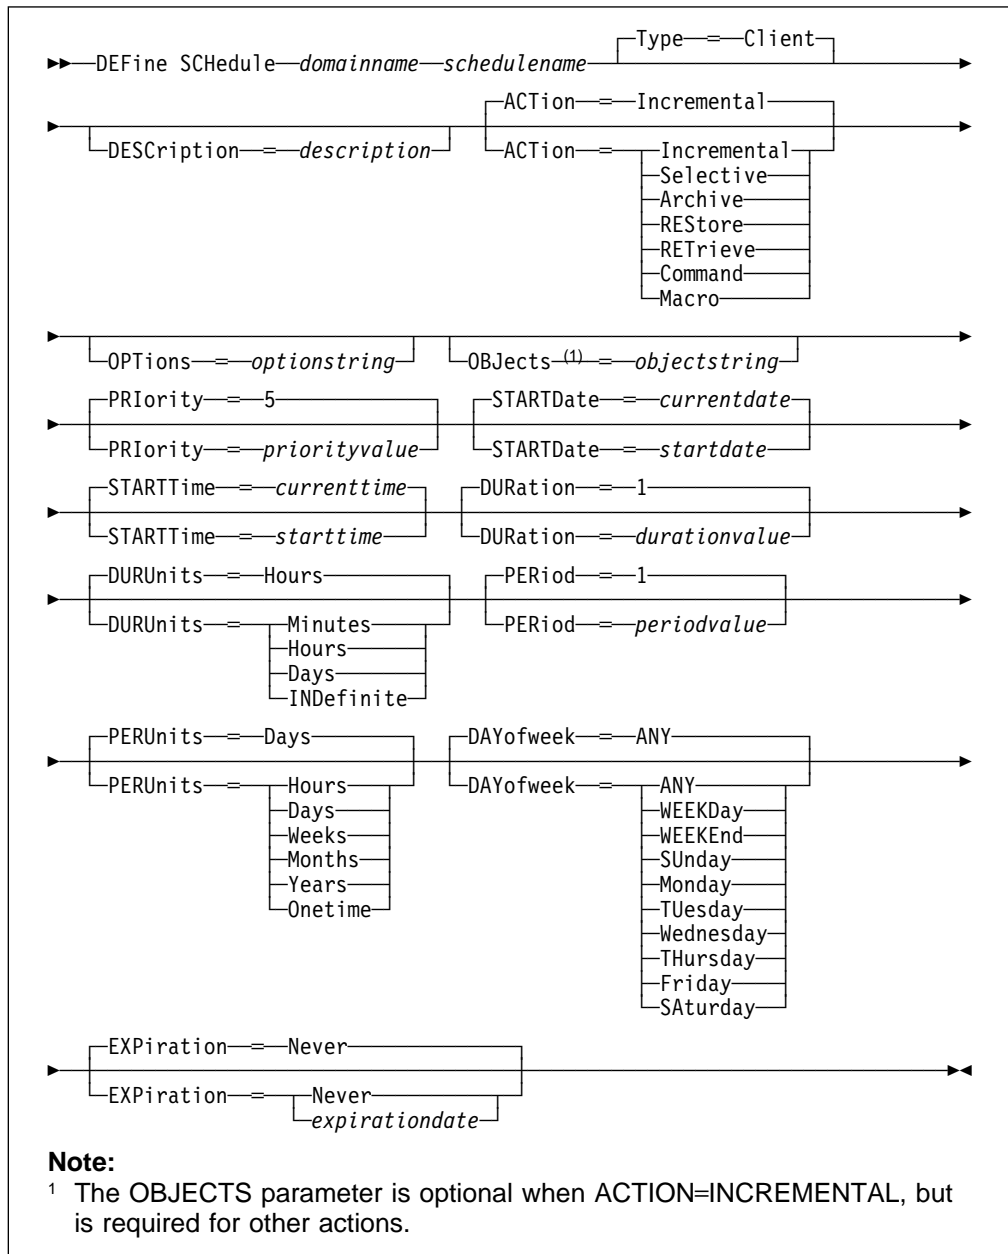
For each schedule, a startup window is specified. The startup window is the time period during which the schedule must be initiated. The schedule will not necessarily complete execution within this window. If the schedule needs to be retried for any reason, the retry attempt must begin before the startup window elapses, or the operation does not restart.

After you define a schedule, use the DEFINE ASSOCIATION command to associate specific clients with the schedule so that they will execute it.

Privilege Class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the schedule belongs.

Syntax for Defining Client Schedules



Parameters for Defining Client Schedules

domainname

Specifies the name of the policy domain to which this schedule belongs. This parameter is required.

schedulename

Specifies the name of the schedule to be defined. This parameter is required. The maximum length of the name is 30 characters.

Type=Client

Specifies that this schedule defines a client schedule. This parameter is optional. The default value is CLIENT. Administrative command schedules are defined by specifying TYPE=ADMINISTRATIVE.

DESCRiption=*description*

Specifies a description of the schedule. This parameter is optional. The maximum length of the description is 255 characters. The description must be enclosed in quotation marks if it contains any blank characters. The default is no description.

ACTion=*actionvalue*

Specifies the action that occurs when this schedule is executed. This parameter is optional. The default value is INCREMENTAL. Possible values are:

Incremental

Specifies that you want this schedule to initiate the backup of any files that have changed since the last backup and for which the specified time between backups has elapsed.

Selective

Specifies that you want this schedule to back up only the files specified with the OBJECTS parameter.

Archive

Specifies that you want this schedule to archive the files specified with the OBJECTS parameter.

REStore

Specifies that you want this schedule to restore the files specified with the OBJECTS parameter.

RETrieve

Specifies that you want this schedule to retrieve the files specified with the OBJECTS parameter.

Command

Specifies that the action is to be a client operating system command string specified with the OBJECTS parameter.

Macro

Specifies that a macro whose file name is specified with the OBJECTS parameter is to be executed.

OPTions=*optionstring*

Specifies options that are supplied to the DSMC command at the time the schedule is executed. This parameter is optional. You can specify most options from the client's option file. For more information, refer to the appropriate client manual. Each option is preceded by a hyphen and delimited with a space. The option string must be enclosed in quotation marks if it contains any blank characters.

When applicable, these options override the options specified by a client node after it has successfully contacted the server.

Do not include the following options because they have no effect on the client's execution of the scheduled command:

MAXCMDRETRIES
 QUERYSCHEDPERIOD
 RETRYPERIOD
 SCHEDLOGNAME
 SCHEDMODE
 TCPCLIENTADDRESS

TCPCLIENTPORT

OBjects=*objectstring*

Specifies the objects for which the specified action is performed. This parameter is required except when ACTION=INCREMENTAL. The allowed objects depend on the action you specify with the ACTION parameter and the operating system used by the clients associated with the schedule.

You can specify more than one object by separating the objects with commas, with no intervening spaces. For example, you can enter a string of file names. If the value for this parameter contains any blank characters, it must be enclosed in quotation marks.

When you specify ACTION=INCREMENTAL and do not specify any objects with this parameter, the DSMC command is invoked without specified objects and attempts to process objects as defined in the client's option file.

PRiority=*priorityvalue*

Specifies the priority value for a schedule. This parameter is optional. You can specify an integer in the range of 1–10, inclusive, with 1 being the highest priority and 10 being the lowest. The default value is 5.

If two or more schedules for a given client have the same window start time, this value determines the order of execution. The schedule with the highest priority starts first. For example, a schedule with PRIORITY=3 starts before a schedule with PRIORITY=5.

STARTDate=*startdate*

Specifies the date for the beginning of the window in which the schedule is first executed. This parameter is optional. The default value is the current date. This parameter is used in conjunction with the STARTTIME parameter to specify when the initial startup window begins.

Possible values are:

MM/DD/YYYY

Specifies the date for the beginning of the startup window in MM/DD/YYYY format (for example, 01/31/1995).

TODAY

Specifies that the beginning of the startup window is on the current date.

TODAY+*numdays* or +*numdays*

Specifies that the beginning of the startup window is a specific number of days after the current date. The maximum value for the number of days is 9999.

For example, to specify that the beginning of the initial startup window is 3 days from now, you can specify STARTDATE=TODAY+3 or STARTDATE=+3.

STARTTime=*starttime*

Specifies the time for the beginning of the window in which the schedule is first executed, in HH:MM format. This parameter is optional. The default value is the current time. This parameter is used in conjunction with the STARTDATE parameter to specify when the initial startup window begins.

DURation=*durationvalue*

Specifies the number of units that define the length of the startup window. This parameter is optional. This value must be in the range of 1–999, inclusive. The default value is 1.

This parameter is used in conjunction with the DURUNITS parameter to specify the length of the startup window. For example, if you specify DURATION=20 and DURUNITS=MINUTES, the schedule must be started within 20 minutes of the start date and start time. The default value for the length of the startup window is 1 hour. The duration of the window must be shorter than the period between windows.

This value is ignored if you specify DURUNITS=INDEFINITE.

DURUnits=*durunitsvalue*

Specifies the time units used to determine the duration of the window in which the schedule can start. This parameter is optional. The default value is HOURS.

This parameter is used in conjunction with the DURATION parameter to specify how long the startup window remains open to execute the schedule. For example, if DURATION=20 and DURUNITS=MINUTES, the schedule must be started within 20 minutes of the start date and start time. The default value for the length of the startup window is 1 hour. The duration of the window must be shorter than the period between windows. Possible values are:

Minutes

Specifies that the duration of the window is defined in minutes.

Hours

Specifies that the duration of the window is defined in hours.

Days

Specifies that the duration of the window is defined in days.

INDefinite

Specifies that the startup window has an indefinite duration. The schedule can run any time after the scheduled start time, until the schedule expires. You cannot specify DURUNITS=INDEFINITE unless you also specify PERUNITS=ONETIME.

PERiod=*periodvalue*

Specifies the length of time between startup windows for this schedule. This parameter is optional. You can specify an integer in the range of 1–999, inclusive. The default value is 1.

This parameter is used in conjunction with the PERUNITS parameter to specify the period between startup windows. For example, if you specify PERIOD=5 and PERUNITS=DAYS (assuming that DAYOFWEEK=ANY), the operation is scheduled every 5 days after the initial start date and start time. The period between startup windows must exceed the duration of each window. The default value for the length of time between startup windows is 1 day.

This value is ignored if you specify PERUNITS=ONETIME.

DEFINE SCHEDULE

PERUnits=*perunitsvalue*

Specifies the time units used to determine the period between startup windows for this schedule. This parameter is optional. The default value is DAYS.

This parameter is used in conjunction with the PERIOD parameter to specify the period between startup windows. For example, if you specify PERIOD=5 and PERUNITS=DAYS (assuming that DAYOFWEEK=ANY), the operation is scheduled every 5 days after the initial start date and start time. The period between startup windows must exceed the duration of each window. The default value for the length of time between startup windows is 1 day. Possible values are:

Hours

Specifies that the time between startup windows for the schedule is in hours.

Days

Specifies that the time between startup windows for the schedule is in days.

Weeks

Specifies that the time between startup windows for the schedule is in weeks.

Months

Specifies that the time between startup windows for the schedule is in months.

Years

Specifies that the time between startup windows for the schedule is in years.

Onetime

Specifies that the schedule executes only once. This value overrides the value specified for the PERIOD parameter.

DAYofweek=*dayofweekvalue*

Specifies the day of the week on which the startup window begins. This parameter is optional. If the beginning of the calculated startup window falls on a day that does not correspond with your value for the DAYOFWEEK parameter, the startup window will be shifted forward in 24-hour increments until the DAYOFWEEK parameter is satisfied.

If you select a value for DAYOFWEEK other than ANY, depending on the values for PERIOD and PERUNITS, schedules may not be executed when you would expect. Use the QUERY EVENT command to project when schedules will be executed and thereby ensure that the desired result is achieved. The default value is ANY. Possible values are:

ANY

Specifies that the startup window can begin on any day of the week.

WEEKDay

Specifies that the startup window can begin on Monday, Tuesday, Wednesday, Thursday, or Friday.

WEEKEnd

Specifies that the startup window can begin on Saturday or Sunday.

SUNday

Specifies that the startup window begins on Sunday.

Monday

Specifies that the startup window begins on Monday.

TUesday

Specifies that the startup window begins on Tuesday.

Wednesday

Specifies that the startup window begins on Wednesday.

THursday

Specifies that the startup window begins on Thursday.

Friday

Specifies that the startup window begins on Friday.

SAturday

Specifies that the startup window begins on Saturday.

EXPIration=*expirationdate*

Specifies the date on which this schedule is no longer used, in MM/DD/YYYY format. This parameter is optional. The default value is NEVER. Possible values are:

Never

Specifies that the schedule never expires.

expirationdate

Specifies the date on which this schedule is no longer used, in MM/DD/YYYY format. If you specify an expiration date, the schedule expires at 23:59:59 on the specified date.

Defining Administrative Command Schedules

Use the DEFINE SCHEDULE command to create a new schedule for the execution of an administrative command.

A schedule determines:

- The administrative command to run
- A startup window that indicates a time range during which the operation can begin
- The time period between startup windows

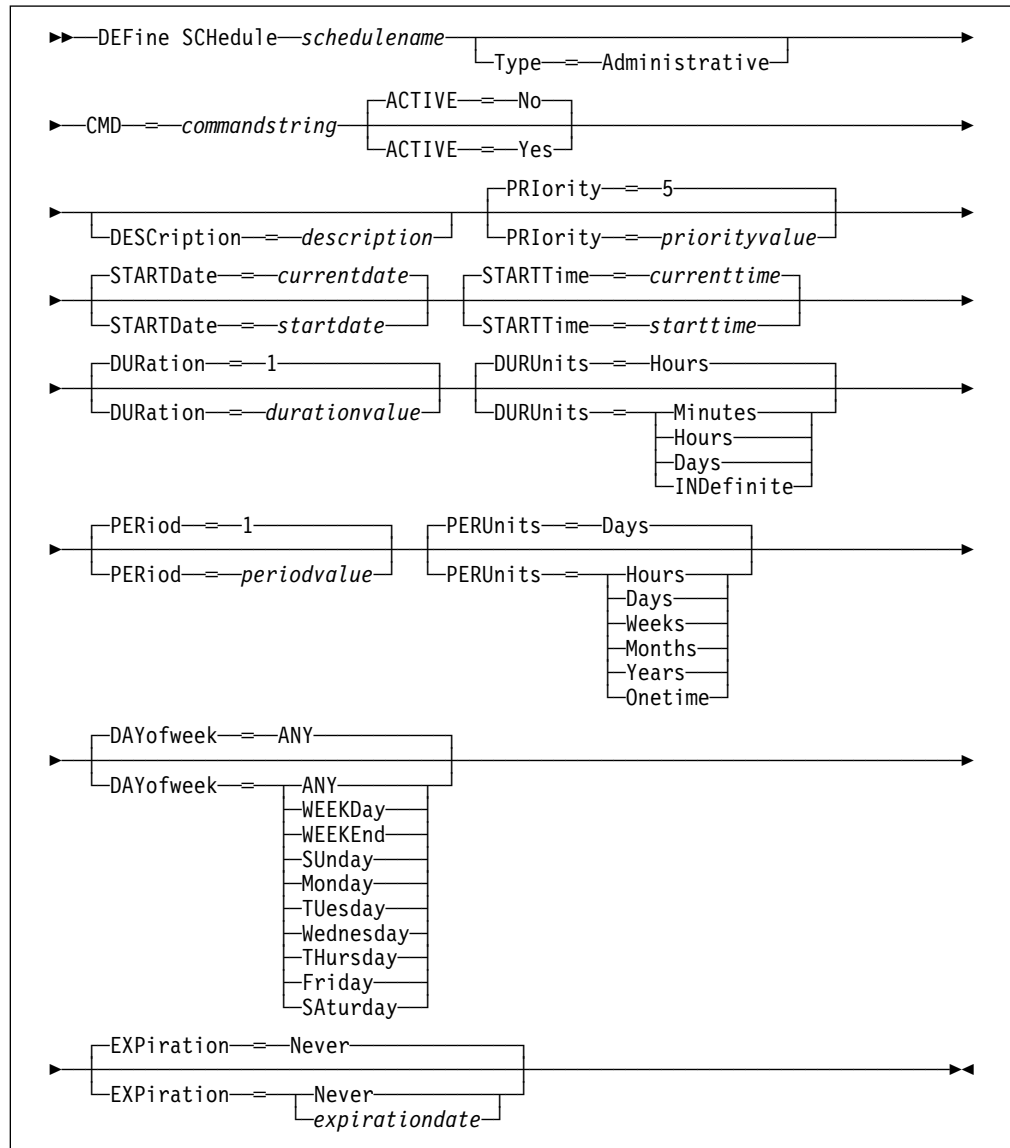
For each schedule, a startup window (or *duration*) is specified. The startup window is the time period during which the schedule must be initiated. The schedule will not necessarily complete execution within this window. If the server is not running when this window starts, but is started before the end of the defined window is reached, the schedule will run when the server is restarted.

DEFINE SCHEDULE

Privilege Class

To issue this command, you must have system privilege.

Syntax for Defining Administrative Command Schedules



Parameters for Defining Administrative Command Schedules

schedulename

Specifies the name of the command schedule to be defined. This parameter is required. The maximum length of the name is 30 characters.

Type=Administrative

Specifies that this schedule defines an administrative command (as opposed to a client command). This parameter is optional. An administrative command is assumed if the CMD parameter is specified.

CMD=commandstring

Specifies the actual administrative command that is to be scheduled. The maximum length of the command is 512 characters. The administrative

command must be enclosed in quotation marks if it contains any blank characters. This parameter is required for administrative command schedules.

ACTIVE=*activevalue*

Specifies if the defined administrative command is immediately eligible for scheduled execution. This parameter is optional. The default value is NO. The administrative command schedule must be set to the active state with the UPDATE SCHEDULE command to be made eligible for scheduled execution. A defined administrative command schedule will not be executed unless it is set to the active state. Possible values are:

YES

Specifies that the defined administrative command is eligible for scheduled execution.

NO

Specifies that the defined administrative command is not eligible for scheduled execution.

DESCRiption=*description*

Specifies a description of the schedule. This parameter is optional. The maximum length of the description is 255 characters. The description must be enclosed in quotation marks if it contains any blank characters. The default is no description.

PRIority=*priorityvalue*

Specifies the priority value for a schedule. This parameter is optional. You can specify an integer in the range of 1–10, inclusive, with 1 being the highest priority and 10 being the lowest. The default value is 5.

If two or more administrative schedules have the same window start time, this value determines the order of execution. The schedule with the highest priority starts first. For example, a schedule with PRIORITY=3 starts before a schedule with PRIORITY=5.

STARTDate=*startdate*

Specifies the date for the beginning of the window in which the schedule is first executed. This parameter is optional. The default value is the current date. This parameter is used in conjunction with the STARTTIME parameter to specify when the initial startup window begins.

Possible values are:

MM/DD/YYYY

Specifies the date for the beginning of the startup window in MM/DD/YYYY format (for example, 01/31/1995).

TODAY

Specifies that the beginning of the startup window is on the current date.

TODAY+*numdays* or +*numdays*

Specifies that the beginning of the startup window is a specific number of days after the current date. The maximum value for the number of days is 9999.

For example, to specify that the beginning of the initial startup window is 3 days from now, you can specify STARTDATE=TODAY+3 or STARTDATE=+3.

DEFINE SCHEDULE

STARTTime=*starttime*

Specifies the time for the beginning of the window in which the schedule is first executed, in HH:MM format. This parameter is optional. The default value is the current time. This parameter is used in conjunction with the STARTDATE parameter to specify when the initial startup window begins.

DURation=*durationvalue*

Specifies the number of units that define the length of the startup window. This parameter is optional. This value must be in the range of 1–999, inclusive. The default value is 1.

This parameter is used in conjunction with the DURUNITS parameter to specify the length of the startup window. For example, if you specify DURATION=20 and DURUNITS=MINUTES, the schedule must be started within 20 minutes of the start date and start time. The default value for the length of the startup window is 1 hour. The duration of the window must be shorter than the period between windows.

This value is ignored if you specify DURUNITS=INDEFINITE.

DURUnits=*durunitsvalue*

Specifies the time units used to determine the duration of the window in which the schedule can start. This parameter is optional. The default value is HOURS.

This parameter is used in conjunction with the DURATION parameter to specify how long the startup window remains open to execute the schedule. For example, if DURATION=20 and DURUNITS=MINUTES, the schedule must be started within 20 minutes of the start date and start time. The default value for the length of the startup window is 1 hour. The duration of the window must be shorter than the period between windows. Possible values are:

Minutes

Specifies that the duration of the window is defined in minutes.

Hours

Specifies that the duration of the window is defined in hours.

Days

Specifies that the duration of the window is defined in days.

INDefinite

Specifies that the startup window has an indefinite duration. The schedule can run any time after the scheduled start time, until the schedule expires. You cannot specify DURUNITS=INDEFINITE unless you also specify PERUNITS=ONETIME.

PERiod=*periodvalue*

Specifies the length of time between startup windows for this schedule. This parameter is optional. You can specify an integer in the range of 1–999, inclusive. The default value is 1.

This parameter is used in conjunction with the PERUNITS parameter to specify the period between startup windows. For example, if you specify PERIOD=5 and PERUNITS=DAYS (assuming that DAYOFWEEK=ANY), the operation is scheduled every 5 days after the initial start date and start time. The period between startup windows must exceed the duration of each

window. The default value for the length of time between startup windows is 1 day.

This value is ignored if you specify PERUNITS=ONETIME.

PERUnits=*perunitsvalue*

Specifies the time units used to determine the period between startup windows for this schedule. This parameter is optional. The default value is DAYS.

This parameter is used in conjunction with the PERIOD parameter to specify the period between startup windows. For example, if you specify PERIOD=5 and PERUNITS=DAYS (assuming that DAYOFWEEK=ANY), the operation is scheduled every 5 days after the initial start date and start time. The period between startup windows must exceed the duration of each window. The default value for the length of time between startup windows is 1 day. Possible values are:

Hours

Specifies that the time between startup windows for the schedule is in hours.

Days

Specifies that the time between startup windows for the schedule is in days.

Weeks

Specifies that the time between startup windows for the schedule is in weeks.

Months

Specifies that the time between startup windows for the schedule is in months.

Years

Specifies that the time between startup windows for the schedule is in years.

Onetime

Specifies that the schedule executes only once. This value overrides the value specified for the PERIOD parameter.

DAYofweek=*dayofweekvalue*

Specifies the day of the week on which the startup window begins. If the beginning of the calculated startup window falls on a day that does not correspond with your value for the DAYOFWEEK parameter, the startup window will be shifted forward in 24-hour increments until the DAYOFWEEK parameter is satisfied.

If you select a value for DAYOFWEEK other than ANY, depending on the values for PERIOD and PERUNITS, schedules may not be executed when you would expect. Use the QUERY EVENT TYPE=ADMINISTRATIVE command to project when schedules will be executed and thereby ensure that the desired result is achieved.

The default value is ANY. Possible values are:

ANY

Specifies that the startup window can begin on any day of the week.

DEFINE SCHEDULE

WEEKDay

Specifies that the startup window can begin on Monday, Tuesday, Wednesday, Thursday, or Friday.

WEEKEnd

Specifies that the startup window can begin on Saturday or Sunday.

SUNday

Specifies that the startup window begins on Sunday.

Monday

Specifies that the startup window begins on Monday.

Tuesday

Specifies that the startup window begins on Tuesday.

Wednesday

Specifies that the startup window begins on Wednesday.

Thursday

Specifies that the startup window begins on Thursday.

Friday

Specifies that the startup window begins on Friday.

SAaturday

Specifies that the startup window begins on Saturday.

EXPIration=expirationdate

Specifies the date on which this schedule is no longer used, in MM/DD/YYYY format. This parameter is optional. The default value is NEVER. Possible values are:

Never

Specifies that the schedule never expires.

expirationdate

Specifies the date on which this schedule is no longer used, in MM/DD/YYYY format. If you specify an expiration date, the schedule expires at 23:59:59 on the specified date.

Examples

Task 1

Define a schedule named MONTHLY_BACKUP that initiates an incremental backup of all associated nodes. Because the specified start date, May 26, 1994 (05/26/1994), is a Thursday and does not match the specified day of the week (Sunday), the initial startup window begins on the first Sunday after May 26, 1994 (05/26/1994). The startup windows for this schedule extend from 01:00 through 03:00. This monthly schedule initiates backup of c: and d: file spaces for all associated nodes.

Command: DEFINE SCHEDULE standard monthly_backup DESCRIPTION="Monthly Backup of c: and d: drives" OBJECTS="c:* d:*" STARTDATE=05/26/1994 STARTTIME=01:00 DURATION=2 DURUNITS=Hours PERIOD=1 PERUNITS=Months DAYOFWEEK=Sunday

Output:

```
ANR2500I Schedule MONTHLY_BACKUP defined in policy domain STANDARD.
```

Task 2

Define a schedule named WEEKLY_BACKUP that initiates an incremental backup of all associated nodes. The initial startup window for this schedule extends from 23:00 on Saturday, July 30, 1994 (07/30/1994), to 03:00 on Sunday, July 31, 1994 (07/31/1994). Subsequent windows begin at 23:00, every Saturday. No messages are returned to the client node when this schedule is executed.

Command: DEFine SChedule os2domain weekly_backup STARTDate=07/30/1994 STARTTime=23:00 DURation=4 DURUnits=Hours PERUnits=Weeks DAYofweek=SAaturday OPTions=-Quiet

Output:

```
ANR2500I Schedule WEEKLY_BACKUP defined in policy domain OS2DOMAIN.
```

Task 3

Define a schedule named BACKUP_ARCHIVEPOOL that backs up the primary storage pool ARCHIVEPOOL to the copy storage pool RECOVERYPOOL. The backup runs at 8 p.m. every two days.

Command: DEFine SChedule BACKUP_ARCHIVEPOOL Type=Administrative CMD="Backup STGpool ARCHIVEPOOL RECOVERYPOOL" ACTIVE=Yes STARTTime=20:00 PERiod=2

Output:

```
ANR2577I Schedule BACKUP_ARCHIVEPOOL defined.
```

Related Commands

Table 41 provides references to related commands.

<i>Table 41 (Page 1 of 2). Commands Related to DEFINE SCHEDULE</i>		
Command	Description	Reference
COPY SCHEDULE	Creates a copy of an existing schedule.	Page 77
DEFINE ASSOCIATION	Associates one or more clients with a schedule.	Page 81
DELETE SCHEDULE	Deletes a schedule from the database.	Page 176
QUERY EVENT	Displays information about scheduled and completed events for selected clients.	Page 290

DEFINE SCHEDULE

Table 41 (Page 2 of 2). Commands Related to DEFINE SCHEDULE

Command	Description	Reference
QUERY SCHEDULE	Displays information about one or more schedules.	Page 335
SET MAXCMDRETRIES	Specifies the maximum number of retries after a failed attempt to execute a scheduled command.	Page 419
SET RETRYPERIOD	Specifies the time between retry attempts by the client scheduler.	Page 430
UPDATE SCHEDULE	Changes the attributes of a schedule.	Page 469

DEFINE STGPOOL (Define a Storage Pool)

Use the DEFINE STGPOOL command to name a set of volumes as a primary storage pool or a copy storage pool. The volumes in a pool belong to the same device class. A primary storage pool is used as the destination for backup files, archive files, or files migrated from client nodes. A copy storage pool is used to store backup copies of files that are in primary storage pools.

During installation, ADSM provides three predefined random access storage pools:

BACKUPPOOL	A storage destination for workstation files that are backed up to the server
ARCHIVEPOOL	A storage destination for workstation files that are archived to the server
SPACEMGPOOL	A storage destination for files migrated from client nodes

Use this command to increase ADSM data storage capacity by defining additional storage pools. You can define storage pools as necessary to manage your storage.

You can restrict administrator authority, by storage pool, to control administrator access to your data storage.

If you want to migrate files from one primary storage pool to another, you must specify the NEXTSTGPOOL parameter. This parameter links primary storage pools together into a hierarchy for data storage.

To determine the name of the device class (and the corresponding device type) to which you want the storage pool assigned, use the QUERY DEVCLASS command.

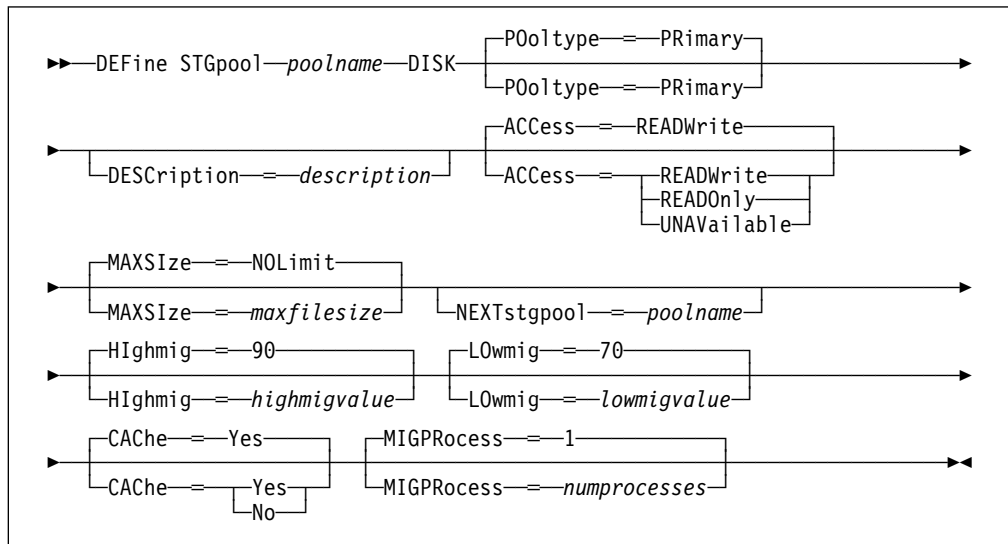
After you define a storage pool, you can use the DEFINE VOLUME command to assign volumes to the storage pool.

The DEFINE STGPOOL command takes three forms, depending on whether you are defining a primary storage pool assigned to random access devices, a primary storage pool assigned to sequential access devices, or a copy storage pool. The syntax and parameters for each form are defined separately.

Privilege Class

To issue this command, you must have system privilege.

Syntax for Defining Random Access Storage Pools



Parameters for Defining a Random Access Storage Pool

poolname

Specifies the name of the storage pool to be defined. This parameter is required. The maximum length of the name is 30 characters.

DISK

Specifies that you want to define a storage pool to the DISK device class. This parameter is required.

POOLTYPE=PRIMARY

Specifies that you want to define a primary storage pool. This parameter is optional. The default value is PRIMARY.

DESCRIPTION=description

Specifies a description of the storage pool to be defined. This parameter is optional. The maximum length of the description is 255 characters. The description must be enclosed in quotation marks if it contains any blank characters. The default is no description.

ACCESS=accessmode

Specifies how users and system processes (like migration and reclamation) can access files in the storage pool. This parameter is optional. The default value is READWRITE. Possible values are:

READWRITE

Specifies that users and system processes (like migration) can read and write to files stored on volumes in the storage pool.

READONLY

Specifies that users can only read files from the volumes in the storage pool.

No new writes are permitted to the volumes in the storage pool from other volumes outside the storage pool. However, system processes (like reclamation) are permitted to move files within the volumes in the storage pool.

If this storage pool has been specified as a subordinate storage pool (with the NEXTSTGPOOL parameter) and is defined as *readonly*, the storage pool is skipped when system processes (like migration) attempt to write files to the storage pool.

UNAVailable

Specifies that users cannot access files stored on volumes in the storage pool.

No new writes are permitted to volumes in the storage pool from other volumes outside the storage pool. However, system processes (like reclamation) are permitted to move files within the volumes in the storage pool. Files can also be moved or copied from this storage pool to another storage pool.

If this storage pool has been specified as a subordinate storage pool (with the NEXTSTGPOOL parameter) and is defined as *unavailable*, the storage pool is skipped when system processes (like migration) attempt to write files to the storage pool.

MAXSize=*maxfilesize*

Specifies the maximum file size for the storage pool. This parameter is optional. The value of this parameter can be either NOLIMIT, or an integer followed by a scale factor. The default value is NOLIMIT, which specifies that there is no maximum size limit for files stored in the storage pool.

To limit the maximum file size, specify an integer (1–999999) followed by a scale factor. Scale factors and their meanings are:

Scale Factor	Meaning
K	Kilobyte
M	Megabyte
G	Gigabyte
T	Terabyte

For example, MAXSIZE=5G specifies that the maximum file size for this storage pool is 5 gigabytes.

If a file exceeds the maximum size and no pool is specified as the NEXTSTGPOOL, the file is not stored. If a file exceeds the maximum size and a pool is specified as the NEXTSTGPOOL, the file is stored in the next storage pool that can accept the file size. If you specify the NEXTSTGPOOL parameter, at least one storage pool in your hierarchy should have a MAXSIZE value of NOLIMIT. By setting MAXSIZE to NOLIMIT for at least one pool, you ensure that no matter what its size, a file can be stored somewhere.

NEXTstgpool=*poolname*

Specifies the name of the storage pool to which files can be migrated. This parameter is optional.

A pool specified with the NEXTSTGPOOL parameter must be a primary storage pool. A copy storage pool cannot be specified.

If you do not specify a value for this parameter, migration cannot occur and files that exceed the value specified for the MAXSIZE parameter cannot be placed in another storage pool.

You cannot create a chain of storage pools that leads to an endless loop through the NEXTSTGPOOL parameter. At least one storage pool in the hierarchy must have no value specified for NEXTSTGPOOL.

Hlghmig=*highmigvalue*

Specifies when ADSM can start migration for this storage pool, as a percentage of the storage pool capacity. This parameter is optional. You can specify an integer from 0–100, inclusive. The default value is 90 percent.

When the storage pool occupancy exceeds the value for this parameter, ADSM can initiate migration, by node, to the next storage pool, as defined with the NEXTSTGPOOL parameter. You can specify HIGHMIG=100 to prevent migration for this storage pool (assuming you have defined the NEXTSTGPOOL parameter).

LOWmig=*lowmigvalue*

Specifies when ADSM can stop migration for this storage pool, as a percentage of the storage pool occupancy. This parameter is optional. You can specify an integer from 0–99. The default value is 70.

When the storage pool occupancy reaches the value for this parameter, ADSM will not initiate migration of another node's files. Because all file spaces that belong to a node are migrated together, the occupancy of the storage pool can fall below the value you specified for this parameter. You can set LOWMIG=0 to permit migration to empty the storage pool.

CAChe=*cachingvalue*

Specifies whether a cached copy of a file that has been migrated remains in the random access storage pool. This parameter is optional. The primary copy of the file is migrated to the next storage pool, as defined with the NEXTSTGPOOL parameter.

Caching improves the retrievability of files that have been migrated to sequential access storage pools by permitting cached copies of files to remain on the storage pool. The cached copies are available, if needed, or can be overwritten when necessary. The default value is YES. Possible values are:

Yes

Specifies that caching is enabled.

No

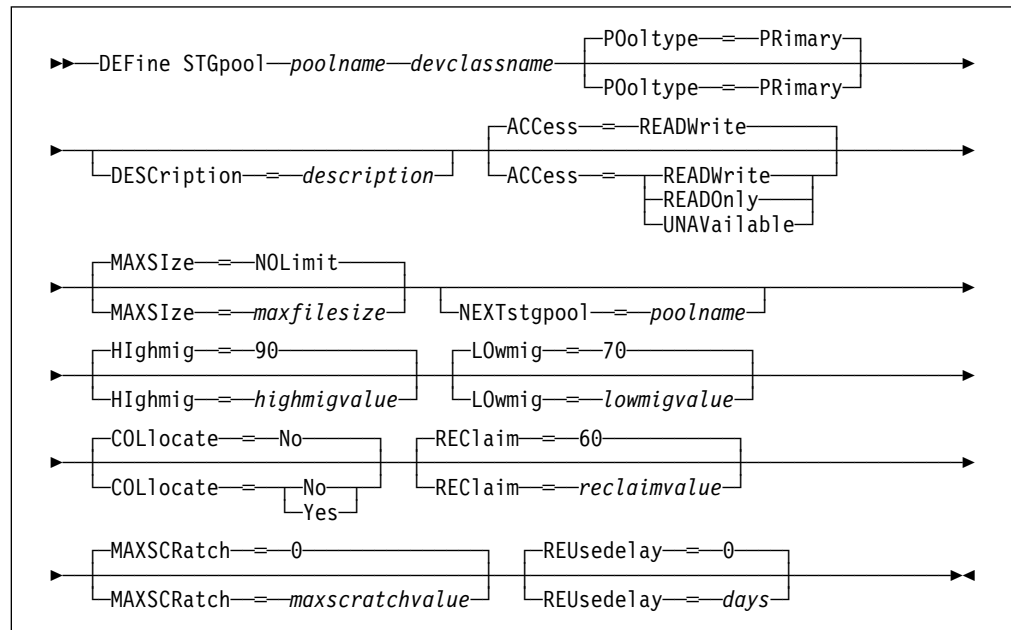
Specifies that caching is disabled.

MIGPRocess=*numprocesses*

Specifies the number of processes that are used for migrating files from this storage pool. This parameter is optional. You can specify an integer from 1–999, inclusive. The default value is 1.

During migration, these processes are performed in parallel to provide the potential for improved migration rates.

Syntax for Defining Primary Sequential Access Storage Pools



Parameters for Defining a Primary Sequential Access Storage Pool

poolname

Specifies the name of the storage pool to be defined. This parameter is required. The maximum length of the name is 30 characters.

devclassname

Specifies the name of the device class to which this storage pool is assigned. This parameter is required. You can specify any device class defined with the DEFINE DEVCLASS command or one of the predefined sequential access device classes (CARTRIDGE or FILE).

POoltype=PRIMARY

Specifies that you want to define a primary storage pool. This parameter is optional. The default value is PRIMARY.

DEscriPtion=*description*

Specifies a description of the storage pool to be defined. This parameter is optional. The maximum length of the description is 255 characters. The description must be enclosed in quotation marks if it contains any blank characters. The default is no description.

ACCess=*accessmode*

Specifies how users and system processes (like migration and reclamation) can access files in the storage pool. This parameter is optional. The default value is READWRITE. Possible values are:

READWrite

Specifies that users and system processes (like migration) can read and write to files stored on volumes in the storage pool.

READOnly

Specifies that users can only read files stored on the volumes in the storage pool.

DEFINE STGPOOL

No new writes are permitted to the volumes in the storage pool from other volumes outside the storage pool. However, system processes (like reclamation) are permitted to move files within the volumes in the storage pool.

If this storage pool has been specified as a subordinate storage pool (with the NEXTSTGPOOL parameter) and is defined as *readonly*, the storage pool is skipped when system processes (like migration) attempt to write files to the storage pool.

UNAVailable

Specifies that users cannot access files stored on volumes in the storage pool.

No new writes are permitted to volumes in the storage pool from other volumes outside the storage pool. However, system processes (like reclamation) are permitted to move files within the volumes in the storage pool. Files can also be moved or copied from this storage pool to another storage pool.

If this storage pool has been specified as a subordinate storage pool (with the NEXTSTGPOOL parameter) and is defined as *unavailable*, the storage pool is skipped when system processes (like migration) attempt to write files to the storage pool.

MAXSize=*maxfilesize*

Specifies the maximum file size for the storage pool. This parameter is optional. The value of this parameter can be either NOLIMIT, or an integer followed by a scale factor. The default value is NOLIMIT, which specifies that there is no maximum size limit for files stored in the storage pool.

To limit the maximum file size, specify an integer (1–999999) followed by a scale factor. Scale factors and their meanings are:

Scale Factor	Meaning
K	Kilobyte
M	Megabyte
G	Gigabyte
T	Terabyte

For example, MAXSIZE=5G specifies that the maximum file size for this storage pool is 5 gigabytes.

If a file exceeds the maximum size and no pool is specified as the NEXTSTGPOOL, the file is not stored. If a file exceeds the maximum size and a pool is specified as the NEXTSTGPOOL, the file is stored in the next storage pool that can accept the file size. If you specify the NEXTSTGPOOL parameter, at least one storage pool in your hierarchy should have a MAXSIZE value of NOLIMIT. By setting MAXSIZE to NOLIMIT for at least one pool, you ensure that no matter what its size, a file can be stored somewhere.

NEXTstgpool=*poolname*

Specifies the name of the storage pool to which files can be migrated. This parameter is optional.

A pool specified with the NEXTSTGPOOL parameter must be a primary storage pool. A copy storage pool cannot be specified.

If you do not specify a value for this parameter, migration cannot occur from this storage pool and files that exceed the value specified for the MAXSIZE parameter cannot be placed in another storage pool.

You cannot create a chain of storage pools that leads to an endless loop through the NEXTSTGPOOL parameter. At least one storage pool in the hierarchy must have no value specified for NEXTSTGPOOL.

Hlghmig=*highmigvalue*

Specifies when ADSM can start migration for this storage pool, as a percentage of the storage pool occupancy. This parameter is optional. You can specify an integer from 1–100, inclusive. The default value is 90 percent.

When the storage pool occupancy exceeds the value for this parameter, ADSM can initiate migration, by volume, to the next storage pool, as defined with the NEXTSTGPOOL parameter. You can specify HIGHMIG=100 to prevent migration for this storage pool (assuming you have defined the NEXTSTGPOOL parameter).

LOWmig=*lowmigvalue*

Specifies when ADSM can stop migration for this storage pool, as a percentage of the storage pool occupancy. This parameter is optional. You can specify an integer from 0–99. The default value is 70.

When the storage pool occupancy reaches the value for this parameter, ADSM will not initiate migration of files from another volume. Because all files that are on a single volume are migrated together, the occupancy of the storage pool can fall below the value you specified for this parameter. You can set LOWMIG=0 to permit migration to empty the storage pool.

COLlocate=*collocation*

Specifies whether ADSM attempts to keep file spaces assigned to as few volumes as possible. This parameter is optional. The default value is NO.

Collocation reduces the number of sequential access media mounts for restore, retrieve, and recall operations. However, collocation increases both the amount of server time needed to collocate files for storing, and the number of volumes required. Collocation may not be useful on systems with only one or two sequential devices. For additional information, see *ADSM Administrator's Guide*.

Possible values are:

Yes

Specifies that collocation is enabled.

No

Specifies that collocation is disabled.

REClaim=*reclaimvalue*

Specifies when to initiate reclamation of fragmented space on volumes in the storage pool, based on the percentage of reclaimable space on a volume. Reclamation makes the fragmented space usable again by moving any remaining active files from one volume to another volume, thus making the original volume available for reuse. This parameter is optional. You can specify an integer from 1–100, inclusive. The default value is 60.

Space becomes fragmented as files on a volume are deleted because of processes such as expiration or file space deletion. The amount of space

DEFINE STGPOOL

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 parameter, the volume is reclaimed. Active files on the volume are rewritten to other volumes in the storage pool, making the original volume available for new files.

It is recommended 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.

MAXSCRatch=*maxscratchvalue*

Specifies the maximum number of scratch volumes that can be requested by ADSM. This parameter is optional. You can specify an integer from 0–100000000, inclusive. The default value is 0.

The value specified for this parameter is used to estimate the total number of volumes available in the storage pool, and the corresponding estimated capacity for the storage pool.

If you do not specify a value for this parameter, no scratch media are requested by ADSM.

Scratch volumes are automatically deleted from the storage pool when they become empty.

REUsedelay=*days*

Specifies the number of days that must elapse after all files have been deleted from a volume, before the volume can be rewritten or returned to the scratch pool. This parameter is optional. You can specify an integer from 0–9999, inclusive. The default value is 0, which means that a volume can be rewritten or returned to the scratch pool as soon as all files have been deleted from the volume.

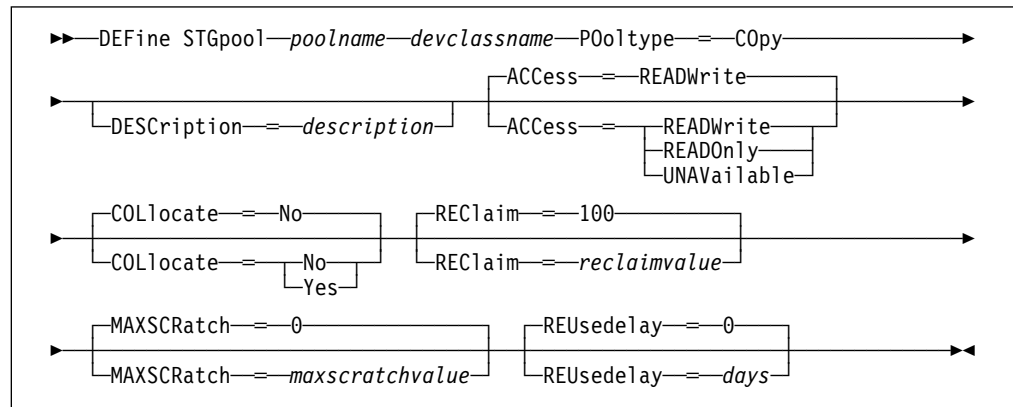
By specifying this parameter, you can ensure that the database could be restored to an earlier level and database references to files in the storage pool would still be valid.

Usage Notes

The migration occupancy represented by the HIGHMIG and LOWMIG parameter is computed differently for random and sequential access storage pools:

- For random access storage pools, the occupancy is computed as a percentage of the total byte capacity of the storage pool.
- For sequential access storage pools, the occupancy is based on the number of volumes in the storage pool that contain files.

Syntax for Defining Copy Storage Pools



Parameters for Defining a Copy Storage Pool

poolname

Specifies the name of the copy storage pool to be defined. This parameter is required. The maximum length of the name is 30 characters.

devclassname

Specifies the name of the sequential access device class to which this copy storage pool is assigned. This parameter is required. You can specify any device class except DISK.

POoltype=COpY

Specifies that you want to define a copy storage pool. This parameter is required.

DESCription=description

Specifies a description of the copy storage pool to be defined. This parameter is optional. The maximum length of the description is 255 characters. The description must be enclosed in quotation marks if it contains any blank characters. The default is no description.

ACCess=accessmode

Specifies how users and system processes (like reclamation) can access files in the copy storage pool. This parameter is optional. The default value is READWRITE. Possible values are:

READWrite

Specifies that files can be read from and written to the volumes in the copy storage pool.

READOnly

Specifies that users can only read files stored on the volumes in the copy storage pool.

No new writes are permitted to the volumes in the copy storage pool from other volumes outside the storage pool. Files cannot be backed up to the copy storage pool, but existing backup copies can be restored. System processes (like reclamation) are permitted to move files within the volumes in the copy storage pool.

UNAVailable

Specifies that users cannot access files stored on volumes in the copy storage pool.

No new writes are permitted to volumes in the copy storage pool from other volumes outside the storage pool. Files cannot be backed up to the copy storage pool, but existing backup copies can be restored.

System processes (like reclamation) are permitted to move files within the volumes in the copy storage pool.

COLlocate=collocation

Specifies whether ADSM attempts to keep file spaces assigned to as few volumes as possible. This parameter is optional. The default value is NO.

Collocation reduces the number of sequential access media mounts for restore, retrieve, and recall operations. However, collocation increases both the amount of server time needed to collocate files for storing, and the number of volumes required. Collocation may not be useful on systems with only one or two sequential devices. For additional information, see *ADSM Administrator's Guide*.

Possible values are:

Yes

Specifies that collocation is enabled.

No

Specifies that collocation is disabled.

REClaim=reclaimvalue

Specifies when to initiate reclamation of fragmented space on volumes in the copy storage pool, based on the percentage of reclaimable space on a volume. Reclamation makes the fragmented space usable again by moving any remaining active files from one volume to another volume, thus making the original volume available for reuse. This parameter is optional. You can specify an integer from 1–100, inclusive. The default value is 100, which means that reclamation is not performed.

The amount of space that can be reclaimed on a volume increases as files on the volume are deleted because of processes such as expiration or file space deletion. It is recommended 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.

When reclamation occurs and how reclamation processing is done depends on whether the volumes are marked as offsite.

For volumes that are not marked as offsite volumes, reclamation can occur only after the volume has been filled, and then begins to empty because of file deletion. When the percentage of reclaimable space on a volume that had once been full rises above the value for the RECLAIM parameter, the volume is reclaimed. Active files on the volume are rewritten to other volumes in the storage pool, making the original volume available for new files.

For offsite volumes, reclamation can occur regardless of whether the volume has ever been filled. An offsite volume is eligible for reclamation when the percentage of unused space on the volume is greater than the reclaim param-

eter value. The unused space includes both space that has never been used on the volume and space that has become empty because of file deletion.

Reclamation processing also works differently for offsite storage pool volumes. When a copy pool volume that is offsite becomes eligible for reclamation, the reclamation process attempts to retrieve the active files on the reclaimable volume from a primary or copy storage pool that is onsite, and then write these files to an available volume in the original copy storage pool. Effectively, these files are moved back to the onsite location, although they could be obtained from the offsite volume after a disaster if a database backup is used which references the files on the offsite volume. Because of the way reclamation works with offsite volumes, use it carefully with copy storage pools.

MAXSCRatch=*maxscratchvalue*

Specifies the maximum number of scratch volumes that can be requested by ADSM. This parameter is optional. You can specify an integer from 0–100000000, inclusive. The default value is 0.

The value specified for this parameter is used to estimate the total number of volumes available in the copy storage pool, and the corresponding estimated capacity for the copy storage pool.

If you do not specify a value for this parameter, no scratch media are requested by ADSM.

Scratch volumes are automatically deleted from the storage pool when they become empty. However, if the access mode for a scratch volume is OFFSITE, the volume will not be deleted from the copy storage pool until the access mode is changed; this allows an administrator to query ADSM for empty, offsite scratch volumes and return these to the onsite location.

REUsedelay=*days*

Specifies the number of days that must elapse after all files have been deleted from a volume, before the volume can be rewritten or returned to the scratch pool. This parameter is optional. You can specify an integer from 0–9999, inclusive. The default value is 0, which means that a volume can be rewritten or returned to the scratch pool as soon as all files have been deleted from the volume.

Important: Specifying this parameter, you can ensure that the database can be restored to an earlier level and database references to files in the copy storage pool are still valid. For this to happen, you must set this parameter to a value greater than the number of days you plan on retaining the oldest database backup. For additional information, refer to *ADSM Administrator's Guide*.

Examples

Task 1

Define a primary storage pool, POOL1, to the DISK device class, with caching enabled. The maximum file size is limited to 5MB, with a high migration threshold of 70 percent, and a low migration threshold of 30 percent. Any files larger than 5MB are stored in subordinate storage pools beginning with the PROG2 storage pool.

DEFINE STGPOOL

Command: DEFine STGpool pool1 disk DESCription= "Disk stgpool, 3390"
MAXSize=5M HIghmig=70 LOwmig=30 CACHE=Yes NEXTstgpool=prog2

Output:

```
ANR2200I Storage pool POOL1 defined (device class DISK).
```

Task 2

Define a primary storage pool named POOLA to the ENDOBYTE device class (with a device type of CARTRIDGE) with a maximum file size of 5MB. Any files larger than 5MB are stored in subordinate pools, beginning with POOL1. Collocation is enabled for the storage pool. As many as 5 scratch volumes can be requested for this storage pool.

Command: DEFine STGpool poola endobyte MAXSize=5M NEXTstgpool=pool1
COLlocate=Yes MAXSCRatch=5

Output:

```
ANR2200I Storage pool POOLA defined (device class  
ENDOBYTE).
```

Task 3

Define a copy storage pool, COPYSTG, to the DC480 device class. Include "Copy Storage Pool" as part of the description for this pool.

Command: DEFine STGpool COPYSTG dc480 pooltype=copy DESCription="Copy
Storage Pool"

Output:

```
ANR2200I Storage pool COPYSTG defined (device class  
DC480).
```

Related Commands

Table 42 provides references to related commands.

Table 42 (Page 1 of 2). Commands Related to DEFINE STGPOOL

Command	Description	Reference
BACKUP DB	Backs up the ADSM database to sequential access volumes.	Page 50
BACKUP STGPOOL	Backs up a primary storage pool to a copy storage pool.	Page 56
DEFINE DEVCLASS	Defines a device class that can be used for storage pools, importing, or exporting.	Page 99

Table 42 (Page 2 of 2). Commands Related to DEFINE STGPOOL

Command	Description	Reference
DEFINE VOLUME	Assigns a volume to be used for storage within a specified storage pool.	Page 148
DELETE STGPOOL	Deletes a storage pool from data storage.	Page 179
QUERY DEVCLASS	Displays information about one or more device classes.	Page 283
QUERY STGPOOL	Displays information about one or more storage pools.	Page 347
RESTORE STGPOOL	Restores files to a primary storage pool from copy storage pools.	Page 391
RESTORE VOLUME	Restores files stored on specified volumes in a primary storage pool from copy storage pools.	Page 396
UPDATE STGPOOL	Changes the attributes of a storage pool.	Page 482

DEFINE VOLUME (Define a Volume in a Storage Pool)

Use the DEFINE VOLUME command to assign a random or sequential access volume to be used for storage within an existing storage pool. You can define a volume to either a primary storage pool or a copy storage pool.

You must define each volume to be used in a storage pool, unless you allow scratch volumes for the storage pool.

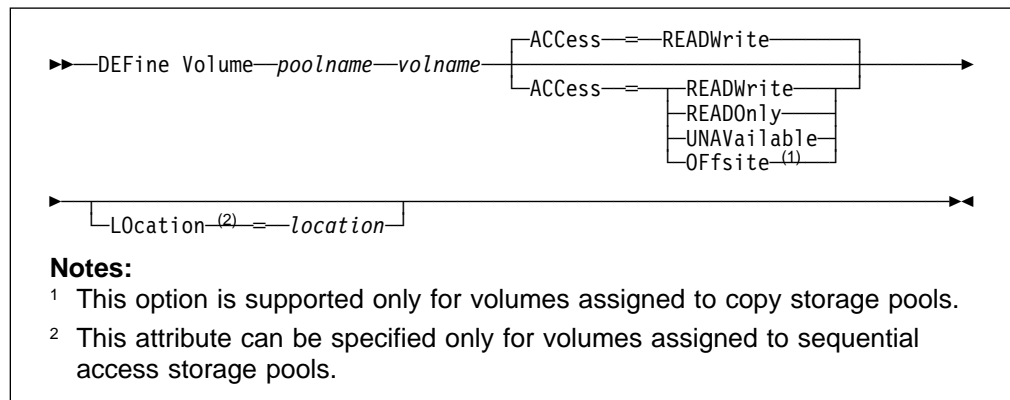
For a random access volume, before issuing this command you must allocate and format the volume by using the ANRFMT job or the ANRFMT2 job found in ADSM.SAMPLIB. For information on using these jobs, see *ADSM Administrator's Guide*.

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

Privilege Class

To issue this command, you must have system privilege, unrestricted storage privilege, or restricted storage privilege for the storage pool to which the volume is assigned.

Syntax



Parameters

poolname

Specifies the name of the storage pool to which the volume is assigned. This parameter is required.

volname

Specifies the name of the storage pool volume to be defined. The volume name that you can specify depends on the device class of the associated storage pool:

- For VSAM linear data sets, specify a standard VSAM cluster name, as defined in the installation JCL. A VSAM cluster name consists of one or more qualifiers, which are delimited by periods. The length of each qual-

ifier is from 1 to 8 alphanumeric characters. The maximum length of the cluster name is 42 characters. An example is ADSM.BFS.POOL2.VOLA.

- For tape cartridges, specify a 1 to 6 alphanumeric tape volume name. An example is DSMT01 (ADSM always converts tape volume names to uppercase).

Any volume name used by ADSM cannot contain any embedded blanks or equal signs.

ACCess=*accessmode*

Specifies how users and system processes (like migration) can access files in the storage pool volume. This parameter is optional. The default value is READWRITE. Possible values are:

READWrite

Specifies that users and system processes (like migration) can read from and write to files stored on the volume.

READOnly

Specifies that users and system processes (like migration) can only read files stored on the volume.

UNAVailable

Specifies that neither users nor system processes (like migration) can access files stored on the volume.

If you define a random access volume as *unavailable*, you cannot vary the volume online. If you define a sequential access volume as *unavailable*, ADSM does not attempt to mount the volume.

OFFsite

Specifies that the volume is at an offsite location from which it cannot be mounted. This option can only be specified for volumes in copy storage pools.

Use this parameter to help you track volumes at offsite locations. ADSM treats volumes designated as offsite differently:

- Mount requests are not generated for volumes assigned to offsite locations
- Data can be reclaimed or moved from offsite volumes by retrieving files from other storage pools
- Empty, offsite scratch volumes are not deleted from the copy storage pool

LOCation=*location*

Specifies the location of the volume. This parameter is optional. It can be specified only for volumes in sequential access storage pools. The location information can be a maximum length of 255 characters. The location must be enclosed in quotation marks if it contains any blank characters. The default is no location.

DEFINE VOLUME

Usage Notes

For sequential access storage pools, it is not necessary to manually define any volumes. You can use the MAXSCRATCH parameter to permit ADSM to mount its own scratch volumes.

Examples

Task 1

A storage pool named STGPOOL1 is assigned to a DISK device class. Define a volume to this storage pool named 901, with READONLY access.

Command: DEFine Volume stgpool1 901 ACCess=READOnly

Output:

```
ANR2206I Volume 901 defined in storage pool STGP00L1
(device class DISK).
```

Task 2

Define a volume named HARVEY.ADSM.STG1 to the POOLA storage pool.

Command: DEFine Volume poola harvey.adsm.stg1

Output:

```
ANR2206I Volume HARVEY.ADSM.STG1 defined in storage pool
POOLA (device class CARTRIDGE).
```

Related Commands

Table 43 provides references to related commands.

Table 43. Commands Related to DEFINE VOLUME

Command	Description	Reference
DELETE VOLUME	Deletes a volume from a storage pool.	Page 185
QUERY VOLUME	Displays information about one or more storage pool volumes.	Page 359
UPDATE VOLUME	Updates the attributes of one or more storage pool volumes.	Page 495

DELETE ASSOCIATION (Delete Node Associations with a Schedule)

Use the DELETE ASSOCIATION command to remove the association between one or more clients and a schedule.

This command does not delete a client from ADSM. It only removes the association of the client from the specified schedule. When you delete the association of a client to a schedule, the client's data is no longer managed by the specified schedule.

If you try to disassociate a client from a schedule to which it is not associated, this command will have no effect for that client.

Privilege Class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the domain to which the schedule belongs.

Syntax

```
►►—DELEte ASSOCiation—domainname—schedulename—nodename—◄◄
```

Parameters

domainname

Specifies the name of the policy domain to which the schedule belongs. This parameter is required.

schedulename

Specifies the name of the schedule from which clients are to be disassociated from. This parameter is required.

nodename

Specifies the name of the node that is to be no longer associated with the schedule. This parameter is required. You can specify a list of clients that are no longer associated with the specified schedule. The items in the list are separated by commas, with no intervening spaces. You can also use a pattern-matching expression to specify a name. All matching clients are disassociated from the specified schedule.

Examples

Task 1

Remove the association between the node named JEFF, assigned to the DOMAIN1 policy domain, and the WEEKLY_BACKUP schedule.

Command: DELEte ASSOCiation domain1 weekly_backup jeff

DELETE ASSOCIATION

Output:

```
Do you wish to proceed? (Y/N)
```

Type Y and press the ENTER key. A message similar to the following is displayed:

```
ANR2511I Node JEFF disassociated from schedule WEEKLY_BACKUP  
in policy domain DOMAIN1.
```

Task 2

Remove the association between selected clients (assigned to the DOMAIN1 policy domain) and the WEEKLY_BACKUP schedule so that this schedule is no longer executed by these clients. The nodes that are disassociated from the schedule contain ABC or XYZ in the node name.

Command: DELEte ASSOCIation domain1 weekly_backup *abc*,*xyz*

Output:

```
Do you wish to proceed? (Y/N)
```

Type Y and press the ENTER key. A message similar to the following is displayed:

```
ANR2511I Node XABCEF disassociated from schedule  
WEEKLY_BACKUP in policy domain DOMAIN1.  
ANR2511I Node ABCONTROL disassociated from schedule  
WEEKLY_BACKUP in policy domain DOMAIN1.  
ANR2511I Node MARTXYZ disassociated from schedule  
WEEKLY_BACKUP in policy domain DOMAIN1.
```

Related Commands

Table 44 provides references to related commands.

Table 44. Commands Related to DELETE ASSOCIATION

Command	Description	Reference
DEFINE ASSOCIATION	Associates one or more clients with a schedule.	Page 81
QUERY ASSOCIATION	Displays the clients associated with one or more schedules.	Page 260

DELETE COPYGROUP (Delete a Copy Group)

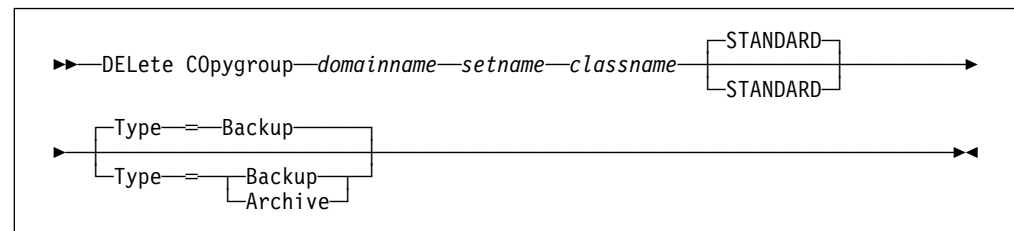
Use the DELETE COPYGROUP command to delete an existing backup or archive copy group from a specified management class. A copy group in the ACTIVE policy set cannot be deleted.

Any files that are bound to a deleted copy group are managed by the default management class (if defined), when the policy set is activated.

Privilege Class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the copy group belongs.

Syntax



Parameters

domainname

Specifies the name of the policy domain to which the copy group to be deleted belongs. This parameter is required.

setname

Specifies the name of policy set to which the copy group to be deleted belongs. This parameter is required.

classname

Specifies the name of the management class to which the copy group to be deleted belongs. This parameter is required.

STANDARD

Specifies the name of the copy group, STANDARD. This parameter is optional. If you do not specify this value, the command defaults to STANDARD.

Type=copygrouptype

Specifies whether a backup or archive copy group is to be deleted. This parameter is optional. The default value is BACKUP. Possible values are:

Backup

Specifies that a backup copy group is to be deleted.

Archive

Specifies that an archive copy group is to be deleted.

DELETE COPYGROUP

Examples

Task 1

Delete the backup copy group, STANDARD, from the ACTIVEFILES management class that is assigned to the VACATION policy set of the EMPLOYEE_RECORDS policy domain.

Command: DELEte C0pygroup employee_records vacation activefiles

Output:

```
Do you wish to proceed? (Y/N)
```

Type Y and press the ENTER key. A message similar to the following is displayed:

```
ANR1531I Backup copy group STANDARD deleted from  
policy domain EMPLOYEE_RECORDS, set VACATION,  
management class ACTIVEFILES.
```

Task 2

Delete the archive copy group, STANDARD, from the MCLASS1 management class that is assigned to the SUMMER policy set of the PROG1 policy domain.

Command: DELEte C0pygroup prog1 summer mclass1 Type=Archive

Output:

```
Do you wish to proceed? (Y/N)
```

Type Y and press the ENTER key. A message similar to the following is displayed:

```
ANR1536I Archive copy group STANDARD deleted from policy  
domain PROG1, set SUMMER, management class MCLASS1.
```

Related Commands

Table 45 provides references to related commands.

Table 45. Commands Related to DELETE COPYGROUP

Command	Description	Reference
DEFINE COPYGROUP	Defines a copy group for backup or archive processing within a specified management class.	Page 83
QUERY COPYGROUP	Displays the attributes of an existing copy group.	Page 270
UPDATE COPYGROUP	Changes one or more attributes of a copy group.	Page 439

DELETE DBBACKUPTRIGGER (Delete the Database Backup Trigger)

Use the DELETE DBBACKUPTRIGGER command to delete the definition for the database backup trigger.

When you delete the definition for the database backup trigger, ADSM does not run backups of the ADSM database automatically and does not remove unnecessary log records from the recovery log. After you delete the database backup trigger, be sure to use the SET LOGMODE command to change the log mode to NORMAL.

Privilege Class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
▶▶—DELEte DBBackuptrigger—————▶▶
```

Parameters

None.

Examples

Task

Delete the definition for the database backup trigger.

Command: DELEte DBBackuptrigger

Output:

```
ANR2286I Database backup trigger deleted.
```

Related Commands

Table 46 provides references to related commands.

Table 46. Commands Related to DELETE DBBACKUPTRIGGER

Command	Description	Reference
DEFINE DBBACKUPTRIGGER	Defines a trigger for automatic ADISM database backups.	Page 91
QUERY DBBACKUPTRIGGER	Displays information about the database backup trigger.	Page 278
SET LOGMODE	Sets the mode for saving records in the recovery log.	Page 416
UPDATE DBBACKUPTRIGGER	Changes one or more attributes of the database backup trigger.	Page 447

DELETE DBVOLUME (Delete a Database Volume)

Use the DELETE DBVOLUME command to delete an existing database volume or database volume copy. This command may require data on the volume to be moved to other database volumes.

A *group of mirrored volumes* can consist of a database volume and up to two volume copies of that database volume. You cannot delete the last synchronized volume in a copy set unless the information on the volume can be moved by ADSM to other synchronized volumes in the database.

This command can generate a background process that can be canceled with the CANCEL PROCESS command. If a DELETE DBVOLUME background process is canceled, some of the database data may have already been moved. To display information on background processes, use the QUERY PROCESS command.

Privilege Class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
▶▶—DELEte DBVolum—volname—————▶▶
```

Parameters

volname

Specifies the name of the database volume to be deleted. This parameter is required.

Examples

Task

Delete the database volume named GLEDDIE.ADSM.PLOG. In the example no data resides on the volume.

Command: DELEte DBVolume gleddie.adsm.plog

Output:

```
ANR2243I Database volume GLEDDIE.ADSM.PLOG deleted.
```

Related Commands

Table 47 provides references to related commands.

Table 47. Commands Related to DELETE DBVOLUME

Command	Description	Reference
CANCEL PROCESS	Cancels a background server process.	Page 64
DEFINE DBCOPY	Creates a volume copy of a database volume.	Page 95
DEFINE DBVOLUME	Assigns a volume to be used for database storage.	Page 97
QUERY DBVOLUME	Displays information about one or more database volumes.	Page 280
QUERY PROCESS	Displays information about one or all background server processes.	Page 332

DELETE DEVCLASS (Delete a Device Class)

Use the DELETE DEVCLASS command to delete an existing device class.

In order to use this command, you must first delete all storage pools assigned to the specified device class and, if necessary, cancel any database backup, export, or import processes that are using the device class.

You cannot delete the device class DISK, which is predefined at installation.

You can delete the device class CARTRIDGE, which is predefined at installation, and any device classes defined by an ADSM administrator.

If you include the DEVCONFIG option in the server options file, the files you specify with that option are automatically updated with the results of this command. When you use this option, the files specified are automatically updated whenever a device class is defined, updated, or deleted.

Privilege Class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
►►—DELeTe DEVcLass—devclassname—————◄◄
```

Parameters

devclassname

Specifies the name of the device class to be deleted. This parameter is required.

Examples

Task

Delete the device class named BITSTOR1. There are no storage pools assigned to the device class.

Command: DELeTe DEVcLass BITSTOR1

Output:

```
ANR2211I Device class BITSTOR1 deleted.
```


Related Commands

Table 48 provides references to related commands.

Table 48. Commands Related to DELETE DEVCLASS

Command	Description	Reference
DEFINE DEVCLASS	Defines a device class that can be used for storage pools, importing, or exporting.	Page 99
QUERY DEVCLASS	Displays information about one or more device classes.	Page 283
UPDATE DEVCLASS	Changes the attributes of a device class.	Page 450

DELETE DOMAIN (Delete a Policy Domain)

Use the DELETE DOMAIN command to delete an existing policy domain. All assigned policy sets, management classes, and copy groups are deleted along with the specified policy domain. A policy domain with registered client nodes cannot be deleted.

Privilege Class

To issue this command you must have system privilege.

Syntax

```
▶▶—DELeTe D0main—domainname————▶▶
```

Parameters

domainname

Specifies the name of the policy domain to be deleted. This parameter is required.

Examples

Task

Delete the EMPLOYEE_RECORDS policy domain.

Command: DELeTe D0main employee_records

Output:

```
Do you wish to proceed? (Y/N)
```

Type Y and press the ENTER key. A message similar to the following is displayed:

```
ANR1501I Policy domain EMPLOYEE_RECORDS deleted.
```

Related Commands

Table 49 provides references to related commands.

Table 49. Commands Related to DELETE DOMAIN

Command	Description	Reference
COPY DOMAIN	Creates a copy of an existing policy domain.	Page 70
DEFINE DOMAIN	Defines a policy domain that can be assigned to one or more clients.	Page 107
QUERY DOMAIN	Displays information about one or more policy domains.	Page 287
UPDATE DOMAIN	Changes the attributes of a policy domain.	Page 458

DELETE EVENT (Delete Event Records)

Use the DELETE EVENT command to delete event records from the database. An event record is created whenever processing of a scheduled command is started or missed.

Event records are deleted automatically based on the value set by the SET EVENTRETENTION command. While it is not necessary for you to manually delete event records, you can use this command to decrease the amount of database space required to hold event records, thus increasing the available capacity of the database.

An event record contains information regarding processing of scheduled work, such as the actual start date and start time, completion date and completion time, status, and results. This information is destroyed if you delete the event record. However, if you delete an event record, it does not alter defined schedules or the work which is pending for associated clients.


This command deletes event records if their scheduled start time is prior to a date and time specified by the administrator, unless the startup window for the event has not yet elapsed. If an event record has a startup window that has not elapsed, the record is not deleted.

This command only deletes the event records that exist at the time the command is executed.

Privilege Class

To issue this command, you must have system privilege or unrestricted policy privilege.

Syntax



Parameters

date

Specifies the date that is used to determine which event records to delete. This parameter is required. This parameter is used with the TIME parameter to specify a date and time for deletion of event records. Any record whose scheduled start occurs prior to the specified date and time is deleted. However, records are not deleted for events whose startup window has not yet passed.

Possible values are:

MM/DD/YYYY

Specifies that you want to delete records created on or before a specific date in MM/DD/YYYY format (for example, 01/31/1995).

TODAY

Specifies that you want to delete records created on or before the current date.

TODAY–numdays or –numdays

Specifies that you want to delete records created on or before the current date minus a specific number of days. The maximum value for the number of days is 9999.

For example, to delete records that are 30 or more days old, you can specify `TODATE=TODAY–30` or `TODATE=–30`.

time

Specifies the time that is used to determine which event records to delete, in HH:MM format. This parameter is optional. The default value is 00:00.

This parameter is used in conjunction with the DATE parameter to specify a time and date for deletion of event records. Any record whose scheduled start occurs prior to the specified time and date is deleted. However, an exception is made for events whose startup window has not yet passed; records are not deleted for such events.

TYPE=typevalue

Specifies if events for client schedules, administrative command schedules, or both, are to be deleted. This parameter is optional. The default value is CLIENT. Possible values are:

Client

Specifies that events for client schedules are to be deleted.

Administrative

Specifies that events for administrative command schedules are to be deleted.

ALI

Specifies that events for both client and administrative command schedules are to be deleted.

Examples

Task 1

Delete records for events with scheduled start times prior to 00:00 on April 15, 1994 (04/15/1994), and whose startup window has passed. Records for these events are deleted regardless of whether the retention period for event records, as specified with the SET EVENTRETENTION command, has passed.

Command: DELeTe EVent 04/15/1994

Output:

Do you wish to proceed? (Y/N)

Type Y and press the ENTER key. A message similar to the following is displayed:

ANR2512I Event records deleted: 18.

DELETE EVENT

Task 2

Delete records for events with scheduled start times prior to 08:00 on May 27, 1994 (05/27/1994), and whose startup window has passed. Records for these events are deleted regardless of whether the retention period for event records, as specified with the SET EVENTRETENTION command, has passed.

Command: DELete EVent 05/27/1994 08:00

Output:

```
Do you wish to proceed? (Y/N)
```

Type Y and press the ENTER key. A message similar to following is displayed:

```
ANR2512I Event records deleted: 36.
```

Related Commands

Table 50 provides references to related commands.

Table 50. Commands Related to DELETE EVENT

Command	Description	Reference
QUERY EVENT	Displays information about scheduled and completed events for selected clients.	Page 290
SET EVENTRETENTION	Specifies the number of days to retain event records.	Page 412

DELETE FILESPACE (Delete Client Node Data from the Server)

Use the DELETE FILESPACE command to delete file spaces from the server. Files that belong to the specified file space are deleted from primary and copy storage pools. Remember that file space names are case sensitive. To determine the correct capitalization for the file space to be deleted, use the QUERY FILESPACE command.

Use the DELETE FILESPACE command to:

- Remove files for users who are not authorized to delete backed up or archived files from storage pools.
- Remove a client node from ADSM. Before removing a client node, you must delete all files belonging to the client node.
- Delete files belonging to a specific owner. When you delete a file space belonging to a specific owner, only the files that have the specified owner name in the file space are deleted.


The DELETE FILESPACE 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 FILESPACE process is canceled or if a system failure occurs, a partial deletion can occur. In this case, a subsequent DELETE FILESPACE command for the same node or owner can delete the remaining data.

This command creates a background process that can be canceled by using the CANCEL PROCESS command. To display information on background processes, use the QUERY PROCESS command. If a DELETE FILESPACE background process is canceled, some or all of the client's file spaces are deleted. If a partial deletion occurred, use the QUERY FILESPACE command to determine which file spaces have not been deleted. From the client node, the remaining files in the file space can be viewed with commands such as QUERY BACKUP. For information about the query commands, refer to the user's guide for your client and platform.

Privilege Class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the client node is assigned.

Syntax



DELETE FILESPACE

Parameters

nodename

Specifies the name of the client node to which the file space belongs. This parameter is required.

filespace name

Specifies the name of the file space to be deleted. This name is case-sensitive. This parameter is required. You can use a pattern-matching expression to specify this name. All matching file space names are included.

Type=datatype

Specifies the type of data to be deleted. This parameter is optional. The default value is ANY. Possible values are:

ANY

Specifies to delete, in order, only backed-up versions of files and archived copies of files.

To delete space-managed files (files migrated from HSM client nodes), you must specify TYPE=SPACEMANAGED.

Backup

Specifies to delete all backup data.

ARchive

Specifies to delete all archive data.

SPacemanaged

Specifies to delete files migrated from a user's local file system (space-managed files from an HSM client). The OWNER parameter is ignored when you specify TYPE=SPACEMANAGED.

OWNer=ownername

Specifies to restrict the data that is deleted from a file space to files belonging to the specified owner. This parameter is optional and is ignored when TYPE=SPACEMANAGED. This parameter only applies to multiuser client systems such as AIX and SunOS.

Examples

Task 1

Delete the file space named C_Drive that belongs to the client node HTANG.

Command: DELEte FIlespace htang C_Drive

Output:

```
Do you wish to proceed? (Y/N)
```

Type Y and press the ENTER key. A message similar to the following is displayed:

```
ANS5104I Process number 3 started.
```


Task 2

Delete the space-managed file space, /MIGTEST, that belongs to client node PEGASUS.

Command: DELEte Filespace pegasus /migtest Type=SPacemanaged

Output:

```
Do you wish to proceed? (Y/N)
```

Type Y and press the ENTER key. A message similar to the following is displayed:

```
ANR0824I Delete Filespace /migtest (space-managed data) for
node PEGASUS started.
ANR0800I DELETE FILESPACE: /migtest for node PEGASUS started
as process 3.
ANR0806I Delete Filespace /migtest complete for node
PEGASUS: 1 files deleted.
```

Related Commands

Table 51 provides references to related commands.

Table 51. Commands Related to DELETE FILESPACE

Command	Description	Reference
CANCEL PROCESS	Cancels a background server process.	Page 64
QUERY ACTLOG	Searches for and displays messages from the server activity log.	Page 252
QUERY FILESPACE	Displays information about data in one or more filespaces that belong to a client.	Page 300
QUERY OCCUPANCY	Displays file space information by storage pool.	Page 324
QUERY PROCESS	Displays information about one or all background server processes.	Page 332
REMOVE NODE	Removes a client from the list of registered nodes for a specific policy domain.	Page 378
RENAME FILESPACE	Renames a client filesystem on the server.	Page 382

DELETE LOGVOLUME (Delete a Recovery Log Volume)

Use the DELETE LOGVOLUME command to delete an existing recovery log volume or recovery log volume copy. This command may require data on the volume to be moved to other database volumes.

A group of mirrored volumes can consist of a recovery log volume and up to two volume copies of that recovery log volume. You cannot delete the last synchronized volume in a copy set unless the information on the volume can be moved by ADSM to other synchronized volumes in the recovery log.

This command can generate a background process that can be canceled with the CANCEL PROCESS command. To display information on background processes, use the QUERY PROCESS command.

Privilege Class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
►►—DELeTe LOGVolumE—volname—————►◄
```

Parameters

volname

Specifies the name of the recovery log volume to be deleted. This parameter is required.

Examples

Task

Delete recovery log volume GLEDDIE.ADSM.PLOG. No data resides on the volume.

Command: DELeTe LOGVolumE gleddie.adsm.plog

Output:

```
ANR2263I Recovery log volume GLEDDIE.ADSM.PLOG deleted.
```

Command: DELEte LOGVolume 300

Output:

```
ANR2264I Delete process initiated for recovery log volume
300 (process id 4).
ANR0245I Deletion complete for recovery log volume.
```

Related Commands

Table 52 provides references to related commands.

Table 52. Commands Related to DELETE LOGVOLUME

Command	Description	Reference
CANCEL PROCESS	Cancels a background server process.	Page 64
DEFINE LOGCOPY	Creates a volume copy of a recovery log volume.	Page 110
DEFINE LOGVOLUME	Defines a disk volume to the recovery log.	Page 112
QUERY LOGVOLUME	Displays information about one or more log volumes.	Page 311
QUERY PROCESS	Displays information about one or all background server processes.	Page 332

DELETE MGMTCLASS (Delete a Management Class)

Use the DELETE MGMTCLASS command to delete an existing management class, including the assigned default management class. You cannot delete a management class in the ACTIVE policy set. All assigned copy groups are deleted along with the specified management class. If you are deleting the default management class, remember a policy set cannot be activated unless it has an assigned default management class.

Privilege Class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the management class belongs.

Syntax

```
►►—DELEte MGmtclass—domainname—setname—classname—————►◄
```

Parameters

domainname

Specifies the name of the policy domain to which the management class belongs. This parameter is required.

setname

Specifies the name of the policy set to which the management class belongs. This parameter is required.

classname

Specifies the name of the management class that you want to delete. This parameter is required.

Examples

Task

Delete the ACTIVEFILES management class from the VACATION policy set of the EMPLOYEE_RECORDS policy domain.

Command: DELEte MGmtclass employee_records vacation activefiles

Output:

```
Do you wish to proceed? (Y/N)
```

Type Y and press the ENTER key. A message similar to the following is displayed:

```
ANR1521I Management class ACTIVEFILES deleted  
from policy domain EMPLOYEE_RECORDS, set VACATION.
```

Related Commands

Table 53 provides references to related commands.

Table 53. Commands Related to DELETE MGMTCLASS

Command	Description	Reference
ASSIGN DEFMGMTCLASS	Assigns an existing management class as the default for a specified policy set.	Page 40
COPY MGMTCLASS	Creates a copy of an existing management class.	Page 73
DEFINE MGMTCLASS	Defines a management class within the policy domain and policy set specified.	Page 114
QUERY MGMTCLASS	Displays information about one or more management classes.	Page 314
UPDATE MGMTCLASS	Changes the attributes of a management class.	Page 460

DELETE POLICYSET (Delete a Policy Set)

Use the DELETE POLICYSET command to delete an existing policy set. You cannot delete the ACTIVE policy set. When a policy set is deleted all management classes and copy groups that belong to the policy set within the specified policy domain are deleted.

Privilege Class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the policy set belongs.

Syntax

```
►►—DELeTe POLiCysEt—domainname—setname—————►►
```

Parameters

domainname

Specifies the name of the policy domain to which the policy set belongs. This parameter is required.

setname

Specifies the name of the policy set to be deleted. This parameter is required.

Examples

Task

Delete the VACATION policy set from the EMPLOYEE_RECORDS policy domain.

Command: DELeTe POLiCysEt employee_records vacation

Output:

```
Do you wish to proceed? (Y/N)
```

Type Y and press the ENTER key. A message similar to the following is displayed:

```
ANR1511I Policy set VACATION deleted from policy domain  
EMPLOYEE_RECORDS.
```

Related Commands

Table 54 provides references to related commands.

Table 54. Commands Related to DELETE POLICYSET

Command	Description	Reference
ACTIVATE POLICYSET	Validates and activates a policy set within a policy domain.	Page 38
COPY POLICYSET	Creates a copy of an existing policy set.	Page 75
DEFINE POLICYSET	Defines a policy set within the specified policy domain.	Page 118
QUERY POLICYSET	Displays information about one or more policy sets.	Page 329
UPDATE POLICYSET	Changes the description of a policy set.	Page 467
VALIDATE POLICYSET	Verifies and reports on conditions the administrator must consider before activating the policy set.	Page 501

DELETE SCHEDULE (Delete a Client or an Administrative Command Schedule)

The DELETE SCHEDULE command takes two forms, depending on whether the schedule applies to client operations or administrative commands. The syntax and parameters for each form are defined separately.

Deleting Client Schedules

Use the DELETE SCHEDULE command to remove one or more existing schedules from the database.

Any client associations to a schedule are removed when the schedule is deleted.

Privilege Class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the specified policy domain.

Syntax for Deleting Client Schedules

```
►►—DELEte SCHedule—domainname—schedulename—[Type=Client]—►►
```

Parameters for Deleting Client Schedules

domainname

Specifies the name of the policy domain to which the schedule belongs. This parameter is required.

schedulename

Specifies the name of the schedule to be deleted. This parameter is required. You can use a pattern-matching expression to specify this name. All matching schedule names are deleted from the specified policy domain.

Type=Client

Specifies that this command should delete a client schedule. This parameter is optional. The default value is CLIENT. Administrative command schedules are deleted by specifying TYPE=ADMINISTRATIVE.

Deleting Administrative Command Schedules

Use the DELETE SCHEDULE command to remove one or more existing schedules from the database.

Privilege Class

To issue this command, you must have system privilege.

Syntax for Deleting Administrative Command Schedules

```
►►—DELeTe SCHedule—schedulename—Type=—Administrative—◄◄
```

Parameters for Deleting Administrative Command Schedules

schedulename

Specifies the name of the schedule to be deleted. This parameter is required. You can use a pattern-matching expression to specify this name. All matching schedule names are deleted.

Type=Administrative

Specifies that the schedule to be deleted is an administrative schedule. This parameter is required for deleting administrative command schedules.

Examples

Task 1

Delete all schedules that belong to the policy domain DOMAIN1.

Command: DELeTe SCHedule domain1 *

Output:

```
Do you wish to proceed? (Y/N)
```

Type Y and press the ENTER key. A message similar to the following is displayed:

```
ANR2501I Schedule WINTER deleted from policy domain DOMAIN1.
ANR2501I Schedule SUMMER deleted from policy domain DOMAIN1.
ANR2501I Schedule SPRING deleted from policy domain DOMAIN1.
```

Task 2

Delete administrative schedule named WEEKLY from the database.

Command: DELeTe SCHedule weekly Type=Administrative

Output:

```
Do you wish to proceed? (Y/N)
```

Type Y and press the ENTER key. A message similar to the following is displayed:

```
ANR2504I Schedule WEEKLY deleted.
```

DELETE SCHEDULE

Related Commands

Table 55 provides references to related commands.

Table 55. Commands Related to DELETE SCHEDULE

Command	Description	Reference
COPY SCHEDULE	Creates a copy of an existing schedule.	Page 77
DEFINE SCHEDULE	Defines a schedule for a client operation or an administrative command.	Page 121
QUERY SCHEDULE	Displays information about one or more schedules.	Page 335
UPDATE SCHEDULE	Changes the attributes of a schedule.	Page 469

DELETE STGPOOL (Delete a Storage Pool)

Use the DELETE STGPOOL command to delete a storage pool.

In order to use this command, you must first delete all volumes assigned to the specified storage pool.

You cannot delete a storage pool that is defined as a subordinate storage pool. For more information on subordinate storage pools, see the NEXTSTGPOOL parameter in the DEFINE STGPOOL command.

Do not delete a storage pool that is specified as a destination for a management class or copy group that is in the ACTIVE policy set.

Privilege Class

To issue this command, you must have system privilege.

Syntax

```
▶▶—DELEte STGpool—poolname————▶▶
```

Parameters

poolname

Specifies the name of the storage pool to be deleted. This parameter is required.

Examples

Task

Delete the storage pool named POOLA.

Command: DELEte STGpool poola

Output:

```
ANR2201I Storage pool POOLA deleted.
```

Related Commands

Table 56 provides references to related commands.

Table 56 (Page 1 of 2). Commands Related to DELETE STGPOOL

Command	Description	Reference
BACKUP STGPOOL	Backs up a primary storage pool to a copy storage pool.	Page 56
DEFINE STGPOOL	Defines a storage pool as a named collection of data storage media.	Page 135

DELETE STGPOOL

Table 56 (Page 2 of 2). Commands Related to DELETE STGPOOL

Command	Description	Reference
QUERY STGPOOL	Displays information about one or more storage pools.	Page 347
UPDATE STGPOOL	Changes the attributes of a storage pool.	Page 482

DELETE VOLHISTORY (Delete Sequential Volume History Information Collected by the Server)

Use the DELETE VOLHISTORY command to delete sequential volume history information collected by the server when the information is no longer needed.

The volume history file contains information about the following types of volumes:

- Volumes used to back up an ADSM database (full or incremental)
- Volumes used to dump an ADSM database
- Volumes used to export administrator, node, policy, or server data
- Storage pool sequential access volumes that have been added
- Storage pool sequential access volumes that have been reused in storage pools through reclamation or MOVE DATA operations
- Storage pool sequential access volumes that have been removed by using the DELETE VOLUME command or removed during reclamation of scratch volumes

To display sequential volume history information collected by the server, use the QUERY VOLHISTORY command.

To save sequential volume history information collected by the server to one or more files, use the BACKUP VOLHISTORY command.

To delete records from volume history information on a regular basis, use the DEFINE SCHEDULE command.

Notes:

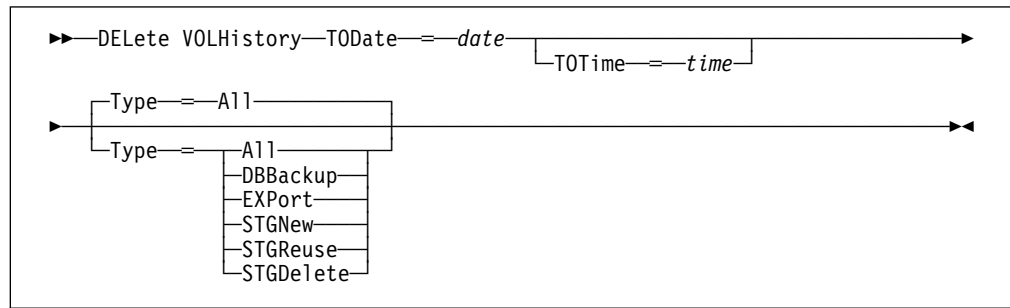
1. Existing volume history files are *not* automatically updated with this command.
2. Do not delete sequential volume history information until you are sure that you no longer need it. For example, do not delete dump volume information or storage volume reuse information unless you have backed up or dumped the database at a later time than that specified for the delete operation.
3. Do not delete the volume history information for database dump, database backup, or export volumes that reside in automated libraries unless you want to return the volumes to scratch status. When the DELETE VOLHISTORY command removes volume information for database dump, database backup, or export volumes, the volumes are automatically returned to scratch status if they reside in automated libraries. The result is that these volumes are available for reuse by the server and the information stored on them may be overwritten when the server reuses the volume for some other purpose, such as storage pool volumes or other database backups.

Privilege Class

To issue this command, you must have system privilege.

DELETE VOLHISTORY

Syntax



Parameters

TODate=*date*

Specifies the date to use to select sequential volume history information to be deleted. ADSM deletes only those records with a date the same as or older than the date you specify.

This parameter is required. Possible values are:

MM/DD/YYYY

Specifies that you want to delete records created on or before a specific date in MM/DD/YYYY format (for example, 01/31/1995).

TODAY

Specifies that you want to delete records created on or before the current date.

TODAY—*numdays* or —*numdays*

Specifies that you want to delete records created on or before the current date minus a specific number of days. The maximum value for the number of days is 9999.

For example, to delete records that are 30 or more days old, you can specify TODATE=TODAY-30 or TODATE=-30.

TOTime=*time*

Specifies the time to use to select records to delete from the volume history file. This parameter is optional.

Of the records that match the date criteria you specify with the TODATE parameter, ADSM deletes only those records with a time that is the same as or earlier than the time you specify with the TOTIME parameter. This parameter is optional. The default value is midnight (00:00:00) on the specified date. Possible values are:

HH:MM:SS

Specifies that you want to delete records created on or before a specific time on the specified date in HH:MM:SS format (for example, 12:30:22).

NOW

Specifies that you want to delete records created on or before the current time on the specified date.

For example, if you issue the DELETE VOLHISTORY command at 9:00 with TOTIME=NOW, ADSM deletes records with a time of 9:00 or earlier on the specified date.

NOW+numhours:numminutes or +numhours:numminutes

Specifies that you want to delete records created on or before the current time plus a specific number of hours (and optionally, minutes) on the specified date.

For example, if you issue the DELETE VOLHISTORY command at 9:00 with TOTIME=NOW+3 or TOTIME=+3, ADSM deletes records with a time of 12:00 or earlier on the specified date.

NOW-numhours:numminutes or -numhours:numminutes

Indicates that you want to delete records created on or before the current time minus a specific number of hours (and optionally, minutes) on the specified date.

For example, if you issue the DELETE VOLHISTORY command at 9:00 with TOTIME=NOW-3:30 or TOTIME=-3:30, ADSM deletes records with a time of 5:30 or earlier on the specified date.

Type=typevalue

Specifies the type of records to delete from the volume history file. This parameter is optional. The default value is ALL. Possible values are:

All

Specifies to delete all records that meet the date and time criteria you specify.

Note: If you previously had an ADSM Version 1 Release 1 or 2 server installed, and you used the DUMP DB command to create online database dumps, records related to the volumes used for those dumps can exist in your volume history information. When you specify TYPE=ALL for the DELETE VOLHISTORY command, ADSM deletes any of those records that meet the date and time criteria you specify.

DBBackup

Specifies to delete only records that both contain information about volumes used for backups (full or incremental), and meet the date and time criteria you specify.

EXPort

Specifies to delete only records that both contain information about volumes used for exporting data, and meet the date and time criteria you specify.

STGNew

Specifies to delete only records that both contain information about new sequential access storage volumes, and meet the date and time criteria you specify.

STGReuse

Specifies to delete only records that both contain information about reused sequential storage pool volumes, and meet the date and time criteria you specify.

STGDelete

Specifies to delete only records that both contain information about deleted sequential storage volumes, and meet the date and time criteria you specify.

DELETE VOLHISTORY

Examples

Task

Delete all volume history information created on or before 03/28/1995.

Command: DELEte VOLHistory T0Date=03/28/1995 Type=A11

Output:

```
Do you wish to proceed? (Yes/No)
```

Type Y and press the Enter key. A message similar to the following is displayed.

```
ANR2467I DELETE VOLHISTORY: 1 sequential volume history  
entries were successfully deleted.
```

Related Commands

Table 57 provides references to related commands.

Table 57. Commands Related to DELETE VOLHISTORY

Command	Description	Reference
BACKUP VOLHISTORY	Records sequential volume history information in external files.	Page 60
DEFINE SCHEDULE	Defines a schedule for a client operation or an administrative command.	Page 121
DELETE VOLUME	Deletes a volume from a storage pool.	Page 185
QUERY VOLHISTORY	Displays sequential volume history information that has been collected by the server.	Page 353

DELETE VOLUME (Delete a Storage Pool Volume)

Use the DELETE VOLUME command to delete a storage pool volume and, optionally, the files within the volume. This command can be used to delete a volume assigned to either a primary or copy storage pool. If during the processing of this command for a primary storage pool volume, ADSM deletes the primary copy of a file (not a cached copy), ADSM also deletes any copies of that file residing in copy storage pools.

To determine the files stored on a volume, use the QUERY CONTENT command.

Before you can delete a storage volume, you must either:

- Use the MOVE DATA command to move all files to another volume.
- Explicitly request to discard all files in the storage volume by specifying DISCARDDATA=YES.

Storage pool volumes may not be totally empty after a move data operation completes if one or more files cannot be relocated to another volume due to input/output errors on the device or if integrity errors were found in the file's structure. In this case, a DELETE operation with DISCARDDATA=YES is required to fully delete the volume. The files with I/O or integrity errors are then lost.

If you specify DISCARDDATA=YES, ADSM creates a server process that deletes files as a series of batch database transactions. If the DELETE VOLUME operation is canceled during processing or if a system failure occurs, a partial deletion can occur. In this case, a subsequent DELETE VOLUME command for the same volume continues deleting the remaining files.

Storage pool volumes cannot be deleted if they are already in use. For example, volumes cannot be deleted if a user is restoring or retrieving a file residing in the storage volume, if ADSM is writing information to the volume, or if reclamation is in process.


This command generates a server process that can be canceled with the CANCEL PROCESS command. To display information on background processes, use the QUERY PROCESS command.

Do not use the DELETE VOLUME command with DISCARDDATA=YES if a restore process (RESTORE STGPOOL or RESTORE VOLUME) is running. The DELETE VOLUME command could cause the restore to be incomplete.

Privilege Class

To issue this command, you must have system privilege, unrestricted storage privilege, or restricted storage privilege for the storage pool to which the volume is defined.

Syntax



DELETE VOLUME

Parameters

volname

Specifies the name of the volume to be deleted. This parameter is required.

DISCARDdata=*discardvalue*

Specifies whether all files stored in the specified volume are deleted from the storage pool. This parameter is optional. No sequential access media mounts are required for this type of deletion. The default value is NO. Possible values are:

No

Specifies that all files stored in the specified volume are not deleted.

Yes

Specifies that all files stored in the specified volume are deleted.

When this option is specified for a primary storage pool volume, the command deletes from the volume both primary copies of files and files that are cached copies (copies left behind by migration processes).

When primary copies of files are deleted from the volume, any copies of these files in copy storage pools are also deleted.

Examples

Task 1

Delete storage pool volume 301 from the storage pool POOL2.

Command: DELEte Volume 301

Output:

```
ANR2220W This command will delete volume 301 from its
storage pool after verifying that the volume contains
no data.
Do you wish to proceed? (Y/N)
```

Type Y and press the ENTER key. A message similar to the following is displayed:

```
ANR2208I Volume 301 deleted from storage pool POOL2.
```

Task 2

Delete storage pool volume HARVEY.ADSM.STG1 from storage pool POOLA. All files on the volume are discarded.

Command: DELEte Volume harvey.adsm.stg1 DISCARDdata=Yes

Output:

```

ANR2220W This command will result in the deletion of all
inventory references to the data on volume
HARVEY.ADSM.STG1, thereby rendering the data unrecoverable.
Do you wish to proceed? (Y/N)

```

Type Y and press the ENTER key. A message similar to the following is displayed:

```

ANR2208I Volume HARVEY.ADSM.STG1 deleted from storage pool
POOLA.

```

Related Commands

Table 58 provides references to related commands.

Table 58. Commands Related to DELETE VOLUME

Command	Description	Reference
CANCEL PROCESS	Cancels a background server process.	Page 64
DEFINE VOLUME	Assigns a volume to be used for storage within a specified storage pool.	Page 148
MOVE DATA	Moves data from a specified storage pool volume to another storage pool volume.	Page 249
QUERY CONTENT	Displays information about one or more files currently residing in a storage pool volume.	Page 265
QUERY PROCESS	Displays information about one or all background server processes.	Page 332
QUERY VOLUME	Displays information about one or more storage pool volumes.	Page 359
UPDATE VOLUME	Updates the attributes of one or more storage pool volumes.	Page 495

DISABLE

DISABLE (Temporarily Prevent Client Node Access to the Server)

Use the DISABLE command to prevent client nodes from starting any new sessions. Current client node sessions are allowed to complete, but no new client node sessions are allowed to start.

This command does not affect system processes (like migration and reclamation).

All registered administrators can access the server when it is disabled with the DISABLE command. To allow client nodes to access the server, use the ENABLE command.

Use the QUERY STATUS command to display the current availability of the server. The DISABLE command can be used before cancelling all client sessions with the CANCEL SESSION command so that the clients cannot automatically reconnect.

Privilege Class

To issue this command, you must have system privilege or operator privilege.

Syntax

```
▶▶—DISAb1e—————▶▶
```

Parameters

None.

Examples

Task

Temporarily prevent new client node sessions from accessing the server.

Command: DISAb1e

Output:

```
ANR2097I Server now disabled for backup/archive client  
access.
```

Related Commands

Table 59 provides references to related commands.

Table 59. Commands Related to DISABLE

Command	Description	Reference
CANCEL SESSION	Cancels one or all active sessions with the server.	Page 66
ENABLE	Resumes server activity following the DISABLE command.	Page 192
QUERY SESSION	Displays information about all active administrator and client sessions with ADSM.	Page 339
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.	Page 342

DISMOUNT DEVICE (Dismount a Volume by Real Device Address)

Use the DISMOUNT DEVICE command to dismount a volume by real device address.

You can only use this command to dismount *idle* volumes. Use the QUERY MOUNT command to determine the status of a mounted volume.

Privilege Class

To issue this command, you must have system privilege or operator privilege.

Syntax

```
▶▶—DISMount Device—deviceaddress————▶▶
```

Parameters

deviceaddress
Specifies the real device address of the volume to be dismounted. This parameter is required.

Examples

Task

Dismount the volume at device address 0FF7.

Command: DISMount Device 0ff7

Output:

```
ANR5280I Command Accepted.
```

Related Commands

Table 60 provides references to related commands.

Table 60. Commands Related to DISMOUNT DEVICE

Command	Description	Reference
DISMOUNT VOLUME	Dismounts a sequential, removable volume by the volume name.	Page 191
QUERY MOUNT	Displays information about mounted sequential access media.	Page 318

DISMOUNT VOLUME (Dismount a Volume by Volume Name)

Use the DISMOUNT VOLUME command to dismount a volume by volume name.

You can only use this command to dismount *idle* volumes. Use the QUERY MOUNT command to determine the status of any mounted volumes.

Privilege Class

To issue this command, you must have system privilege or operator privilege.

Syntax

```
►►—DISMount Volume—volname—◄◄
```

Parameters

volname

Specifies the name of the volume to be dismounted. This parameter is required.

Examples

Task

Dismount the volume BTV005.

Command: DISMount Volume btv005

Output:

```
ANR5280I Command Accepted.
```

Related Commands

Table 61 provides a reference to a related command.

Table 61. Command Related to DISMOUNT VOLUME

Command	Description	Reference
DISMOUNT DEVICE	Dismounts a sequential, removable volume by the real device address.	Page 190
QUERY MOUNT	Displays information about mounted sequential access media.	Page 318

ENABLE (Resume User Activity on the Server)

Use the ENABLE command to permit client node activity to begin after issuing the DISABLE command.

This command does not affect system processes (like migration and reclamation).

Use the QUERY STATUS command to display the current availability of the server.

Privilege Class

To issue this command, you must have system privilege or operator privilege.

Syntax

```
▶▶—ENable—◀◀
```

Parameters

None.

Examples

Task

Resume normal operation, permitting client nodes to access the server.

Command: ENable

Output:

```
ANR2096I Server now enabled for backup/archive access.
```

Related Commands

Table 62 provides references to related commands.

<i>Table 62. Commands Related to ENABLE</i>		
Command	Description	Reference
CANCEL SESSION	Cancels one or all active sessions with the server.	Page 66
DISABLE	Prevents new clients from accessing ADSM but permits existing sessions to continue.	Page 188
QUERY SESSION	Displays information about all active administrator and client sessions with ADSM.	Page 339
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.	Page 342

EXPIRE INVENTORY (Manually Start Inventory Expiration Processing)

Use the EXPIRE INVENTORY command to manually start inventory expiration processing. Only one expiration process is allowed at any time; therefore, if expiration processing is currently running, another process cannot be started. Expiration eligibility is determined by policy information assigned to either the backup or archive files. If a file is eligible for expiration, the storage used by this file is deallocated.

Automatic expiration processing is controlled by the EXPINTERVAL parameter specified in the ADSM server options file. If EXPINTERVAL is set to 0, automatic expiration does not occur and issuing this command is necessary to start expiration processing. For additional information about the EXPINTERVAL parameter and the ADSM server options file, refer to *ADSM Installing the Server and Administrative Client*.

This command creates a background process that can be canceled with the CANCEL PROCESS command. To display information on background processes, use the QUERY PROCESS command.

The inventory expiration process removes client backup versions and archive file copies from server storage under the terms specified by the management classes to which the files are bound. The management class copygroup parameters, RETEXTRA and RETVER, control backup version and archive copy deletion, respectively. Unexpired backup versions and archive copies are considered by the license audit process in determining the amount of storage managed by the server.

Privilege Class

To issue this command, you must have system privilege.

Syntax

```
▶▶—EXPIre Inventory—————▶▶
```

Parameters

None.

Examples

Task

Issue the EXPIRE INVENTORY command.

Command: EXPIre Inventory

Output:

```
ANR0811I Inventory client file expiration started as
process 2.
ANR0812I Inventory file expiration process 2 completed:
deleted 25 backup files and 3 archive files.
```

Related Commands

Table 63 provides references to related commands.

Table 63. Commands Related to EXPIRE INVENTORY

Command	Description	Reference
AUDIT LICENSES	Checks for compliance with defined licenses.	Page 43
CANCEL PROCESS	Cancels a background server process.	Page 64
DEFINE COPYGROUP	Defines a copy group for backup or archive processing within a specified management class.	Page 83
QUERY PROCESS	Displays information about one or all background server processes.	Page 332

EXPORT ADMIN (Copy Administrator Information to Sequential Media)

Use the EXPORT ADMIN command to export administrator and authority definitions for one or more administrators from the ADSM server to sequential media.

The EXPORT ADMIN command exports administrator information such as:

- Administrator name, password, and contact information
- Administrative privilege classes the administrator has been granted

The EXPORT ADMIN command allows administrative information to be copied to another server. In order to copy ADSM data and then restore it to a new server, the device class must be supported on both server platforms.

ADSM client data can be moved between ADSM servers via export and import processing if the same removable media type is supported on both platforms.

This command is useful for permitting later retrieval of the server information for copying the information to a new server.

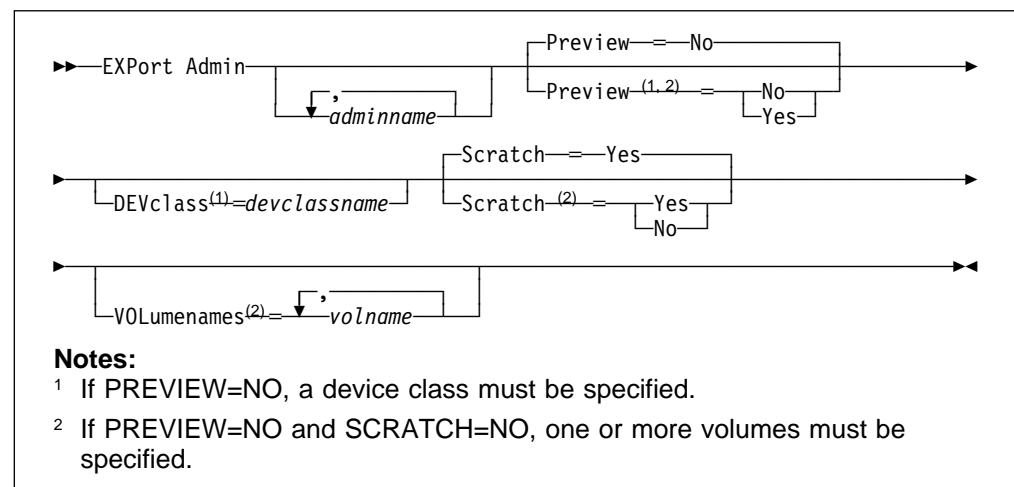
You can use the QUERY ACTLOG command to view the actual status information indicating the size and the success or failure of the export operation. You can also view this information from the server console.

This command generates a background process that can be canceled with the CANCEL PROCESS command. If an EXPORT ADMIN background process is canceled, the sequential access volumes holding the exported data are incomplete and should not be used for importing data. To display information on background processes, use the QUERY PROCESS command.

Privilege Class

To issue this command, you must have system privilege.

Syntax



Parameters

adminname

Specifies a list of administrators for which administrator information is to be exported. This parameter is optional. The items in the list are separated by commas, with no intervening spaces. You can use pattern-matching expressions to specify names. All matching administrators are included in the list. If you do not specify a value for this parameter, information related to all administrators is exported.

Preview=*value*

Specifies whether you want to preview the results of the export operation, without actually exporting administrator information. You can use this parameter to preview how many bytes of data are transferred, allowing you to determine how many volumes will be required to hold the data. This parameter is optional. The default value is NO. Possible values are:

Yes

Specifies that you want to preview the results of the export operation, without exporting administrator information. Information on the number and types of objects exported, together with the number of bytes transferred are reported to the server console and the activity log. If you specify this value, you do not need to specify device or volume parameters.

No

Specifies that the administrator information is to be exported. If you specify this value, you must specify a value for the DEVCLASS parameter.

DEVclass=*devclassname*

Specifies the device class to which the export data is to be written. This parameter is required if you specified PREVIEW=NO. You cannot specify the DISK device class.

If the mount limit for the device class specified has already been reached (all drives are busy) when the export runs, ADSM automatically cancels lower priority operations, such as reclamation, to make a mount point available for the export.

Scratch=*scratchvalue*

Specifies whether scratch volumes can be used for export. The default value is YES. Possible values are:

Yes

Specifies that scratch volumes can be used for export. If you include a list of volumes on the VOLUMENAMES parameter, then scratch volumes are used only if there is not enough space to store the exported files on the volumes specified. If the device type associated with the specified device class is FILE, file names for scratch volumes are generated based on a time stamp.

No

Specifies that scratch volumes cannot be used for export. You must include a list of volumes on the VOLUMENAMES parameter to contain all of the export data (unless PREVIEW=YES). To determine how many

volumes you may need, you can run the command specifying PREVIEW=YES.

VOLumentnames=volname

Specifies one or more sequential access volumes that are to be used to contain exported files. This parameter is optional. The items in the list are separated by commas, with no intervening spaces. Any volume name used by ADSM cannot contain any embedded blanks or equal signs. A volume name must be specified if SCRATCH=NO and PREVIEW=NO. If you do not specify a volume and SCRATCH=YES, scratch volumes are used. Use the following naming conventions when specifying volumes for each device type.

CARTRIDGE

Specify 1–6 alphanumeric characters, for example DSMT01.

FILE

Specify any supported file name, for example /exndata/mt1.

Examples

Task

Export the administrator information for all defined administrators to 3480 cartridge tape (assuming a device class named SQUARE1).

The number and type of objects exported are reported to the system console and in the activity log.

Command: EXPort Admin DEVclass=square1 Scratch=Yes

Output:

```
ANR0609I EXPORT ADMIN started as process 9.
ANS5104I Process number 9 started.
```

Related Commands

Table 64 provides references to related commands.

Table 64 (Page 1 of 2). Commands Related to EXPORT ADMIN

Command	Description	Reference
CANCEL PROCESS	Cancels a background server process.	Page 64
EXPORT NODE	Copies client node information to external media.	Page 199
EXPORT POLICY	Copies policy information to external media.	Page 204
EXPORT SERVER	Copies all or part of the server to external media.	Page 208
IMPORT ADMIN	Restores administrative information from external media.	Page 227

Table 64 (Page 2 of 2). Commands Related to EXPORT ADMIN

Command	Description	Reference
QUERY ACTLOG	Searches for and displays messages from the server activity log.	Page 252
QUERY PROCESS	Displays information about one or all background server processes.	Page 332

EXPORT NODE (Copy Client Node Information to Sequential Media)

Use the EXPORT NODE command to export client node definitions from the ADSM server to sequential media. Each client node definition includes:

- User ID, password, and contact information
- Name of the policy domain to which the client is assigned
- File compression status
- Whether the user has the authority to delete backed up or archived files from data storage
- Whether the client node ID was locked at the time of export

Optionally, the EXPORT NODE command exports node information such as:

- File space definitions
- Data from files migrated from HSM client nodes
- Archive or backup file copy data
- Access authorization information pertaining to the file spaces reported

The EXPORT NODE command allows client node information to be copied to another server. In order to backup ADSM data and then restore it to a new server, the sequential access volume device must be supported on both platforms.

ADSM client data can be moved between ADSM servers via export and import processing if the same removable media type is supported on both platforms.

If a file from a primary storage pool cannot be exported because it is damaged or because it resides on a volume marked as UNAVAILABLE or OFFLINE, ADSM attempts to obtain a copy of the file from a copy storage pool.

This command is useful for permitting later retrieval of server information for copying the information to a new server.

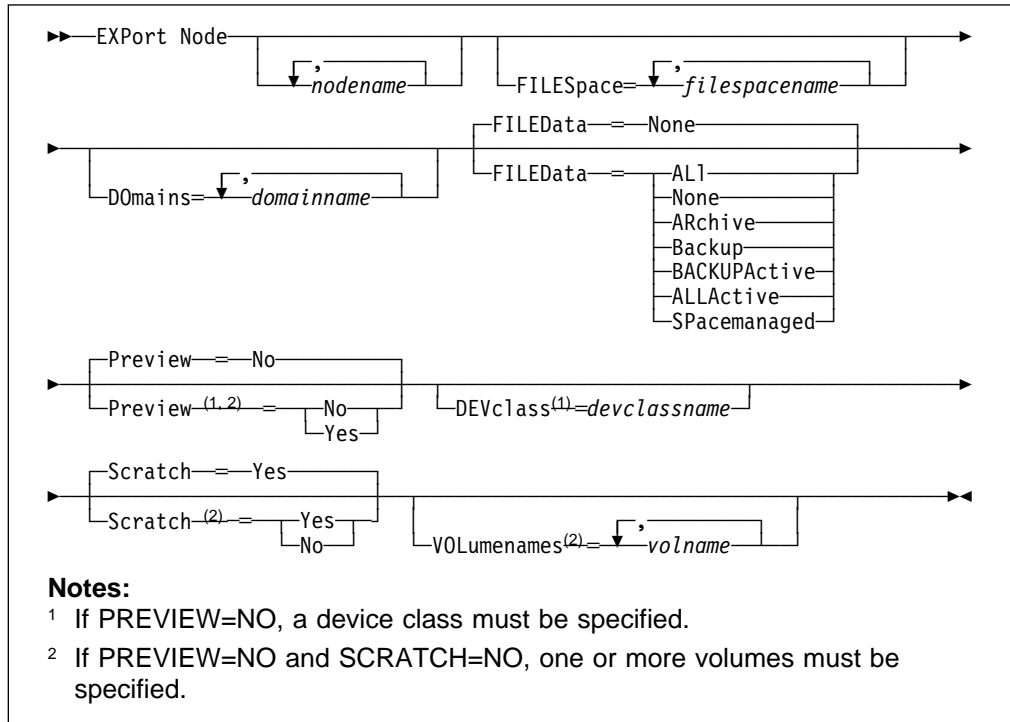
You can use the QUERY ACTLOG command to view the actual status information indicating the size and the success or failure of the export operation. You can also view this information from the server console.

This command generates a background process that can be canceled with the CANCEL PROCESS command. If an EXPORT NODE background process is canceled, the sequential volume media holding the exported data is incomplete and should not be used for importing data. To display information on background processes, use the QUERY PROCESS command.

Privilege Class

To issue this command, you must have system privilege.

Syntax



Parameters

nodename

Specifies a list of client node names for which information is to be exported. This parameter is optional. You can specify a list of node names for which node information is to be exported. The items in the list are separated by commas, with no intervening spaces. You can use pattern-matching expressions to specify names. All matching nodes are included in the list. If you do not specify a value for this parameter, information related to all nodes is exported.

FILESpace=filespacename

Specifies a list of file space names for which data is to be exported. This parameter is optional. You can specify a list of file space names for which data information is to be exported. The items in the list are separated by commas, with no intervening spaces. You can use a pattern-matching expression to specify a name. All matching names are included in the list. If you do not specify a value for this parameter, information related to all files space names is exported.

DOmains=domainname

Specifies a list of policy domain names from which nodes should be exported. This parameter is optional. The items in the list are separated by commas, with no intervening spaces. For a node to be exported it must belong to one of the specified domains. You can use a pattern-matching expression to specify a name. All matching domains are included in the list. If you do not specify a value for this parameter, information related to all domains is exported.

FILEData=*filedatatype*

Specifies the type of files that should be exported for all nodes defined to the server. This parameter is optional. The default value is NONE.

The device class used by the file data is determined by the device class for the storage pool. If it is the same device class specified in this command, then two drives are needed to export node information. The mount limit value for the device class needs to be set to at least 2.

The following descriptions mention *active* and *inactive* file copies. An active file copy is the most recent backup copy for a file that still exists on the client workstation. All other file copies are called inactive copies. Possible values are:

ALI

Specifies that ADSM exports all backup version of files, all archived files, and all files migrated from a user's local file system.

None

Specifies that no files are to be exported; only node definitions are to be exported.

ARchive

Specifies that only archived files are to be exported.

Backup

Specifies that only backup versions, whether active or inactive, are to be exported.

BACKUPActive

Specifies that only active backup versions are to be exported.

ALLActive

Specifies that ADSM exports all active backup versions of files, all archived files, and all files migrated from a user's local file system.

SPacemanaged

Specifies that only files which have been migrated from a user's local file system are to be exported.

Preview=*value*

Specifies whether you want to preview the results of the export operation, without actually exporting information. You can use this parameter to preview how many bytes of data are transferred, allowing you to determine how many tapes will be required to hold the data. This parameter is optional. The default value is NO. Possible values are:

Yes

Specifies that you want to preview the results of the export operation, without actually exporting node information. Information on the number and types of objects exported, together with the number of bytes transferred are reported to the server console and the activity log. If you specify this value, you do not need to specify device or volume parameters.

No

Specifies that the node information is to be exported. If you specify this value, you must specify a value for the DEVCLASS parameter.

DEVclass=*devclassname*

Specifies the device class to which export data is to be written. This parameter is required if PREVIEW=NO. You cannot specify the DISK device class.

If the mount limit for the device class specified has already been reached (all drives are busy) when the export runs, ADSM automatically cancels lower priority operations, such as reclamation, to make a mount point available for the export.

Scratch=*scratchvalue*

Specifies whether scratch volumes can be used for export. The default is SCRATCH=YES. Possible values are:

Yes

Specifies that scratch volumes are to be used for export. If you include a list of volumes on the VOLUMENAMES parameter, then scratch volumes are used only if there is not enough space to store the exported files on the volumes specified. If the device type associated with the specified device class is FILE, file names for scratch volumes are generated based on a time stamp.

No

Specifies that scratch volumes cannot be used for export. You must include a list of volumes on the VOLUMENAMES parameter to contain all of the export data, unless PREVIEW=YES. To determine how many volumes you may need, specify PREVIEW=YES.

VOLUmenames=*volname*

Specifies one or more sequential access volumes that are to be used to contain exported files. This parameter is optional. The items in the list are separated by commas, with no intervening spaces. Any volume name used by ADSM cannot contain any embedded blanks or equal signs. A volume name must be specified if SCRATCH=NO and PREVIEW=NO. If you do not specify a volume and SCRATCH=YES, scratch volumes are used. Use the following naming conventions when specifying volumes for each device type.

CARTRIDGE

Specify 1–6 alphanumeric characters, for example DSMT01.

FILE

Specify any supported file name, for example /exndata/mt1.

Examples

Task

Export the node definitions and active backup data for all nodes defined in the STANDARD policy domain to tape (assuming a device class of MENU2).

The number and type of objects exported are reported to the system console and in the activity log.

Command: EXPort Node D0mains=Standard DEVclass=menu2 Scratch=Yes
FILEData=BACKUPActive

Output:

```

ANR0609I EXPORT NODE started as Process 12.
ANS5104I Process number 12 started.

```

Related Commands

Table 65 provides references to related commands.

Table 65. Commands Related to EXPORT NODE

Command	Description	Reference
CANCEL PROCESS	Cancels a background server process.	Page 64
EXPORT ADMIN	Copies administrative information to external media.	Page 195
EXPORT POLICY	Copies policy information to external media.	Page 204
EXPORT SERVER	Copies all or part of the server to external media.	Page 208
IMPORT NODE	Restores client node information from external media.	Page 230
QUERY ACTLOG	Searches for and displays messages from the server activity log.	Page 252
QUERY PROCESS	Displays information about one or all background server processes.	Page 332

EXPORT POLICY (Copy Policy Information to Sequential Media)

Use the EXPORT POLICY command to export policy information for one or more policy domains from the ADSM server to sequential media.

The EXPORT POLICY command exports policy information and client node data such as:

- Policy domain definitions
- Policy set definitions, including the active policy set
- Management class definitions, including the default management class
- Backup copy group and archive copy group definitions
- Schedule definitions and client node associations

The EXPORT POLICY command allows all policy information to be copied to another server. In order to backup ADSM data and then restore it to a new server, the sequential access device must be supported on both server platforms.

ADSM client data can be moved between ADSM servers via export and import processing if the same removable media type is supported on both platforms.

This command is useful for permitting later retrieval of the server information for copying the information to a new server.

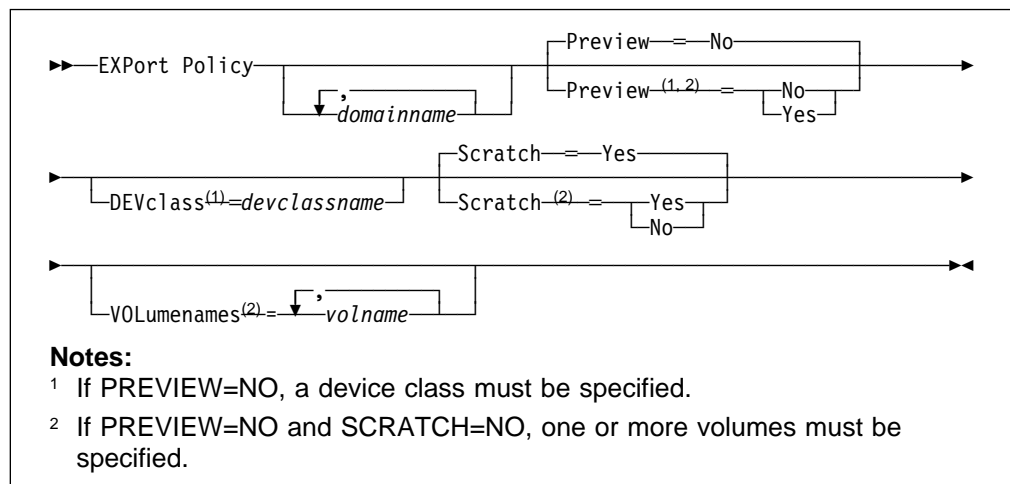
You can use the QUERY ACTLOG command to view the actual status information indicating the size and the success or failure of the export operation. You can also view this information from the server console.

This command generates a background process that can be canceled with the CANCEL PROCESS command. If an EXPORT POLICY background process is canceled, the sequential access media holding the exported data is incomplete and should not be used for importing data. To display information on background processes, use the QUERY PROCESS command.

Privilege Class

To issue this command, you must have system privilege.

Syntax



Parameters

domainname

Specifies a list of policy domains for which policy information is to be exported. This parameter is optional. You can specify a list of policy domains for which policy information is to be exported. The items in the list are separated by commas, with no intervening spaces. You can use pattern-matching expressions to specify names. All matching policy domains are included in the list. If you do not specify a value for this parameter, information related to all policy domains is exported.

Preview=*value*

Specifies whether you want to preview the results of the export operation, without actually exporting information. You can use this parameter to preview how many bytes of data are transferred, allowing you to determine how many volumes will be required to hold the data. This parameter is optional. The default value is NO. Possible values are:

Yes

Specifies that you want to preview the results of the export operation, without actually exporting policy information. Information on the number and types of objects exported, together with the number of bytes transferred are reported to the server console and the activity log. If you specify this value, you do not need to specify device or volume parameters.

No

Specifies that the policy information is to be exported. If you specify this value, you must specify a value for the DEVCLASS parameter.

DEVclass=*devclassname*

Specifies the device class to which export data is to be written. This parameter is required if PREVIEW=NO. You cannot specify the DISK device class.

If the mount limit for the device class specified has already been reached (all drives are busy) when the export runs, ADSM automatically cancels lower priority operations, such as reclamation, to make a mount point available for the export.

Scratch=*scratchvalue*

Specifies whether scratch volumes can be used for export. The default is SCRATCH=YES. Possible values are:

Yes

Specifies that scratch volumes are to be used for export. If you include a list of volumes on the VOLUMENAMES parameter, then scratch volumes are used only if there is not enough space to store the exported files on the volumes specified. If the device type associated with the specified device class is FILE, file names for scratch volumes are generated based on a time stamp.

No

Specifies that scratch volumes cannot be used for export. You must include a list of volumes on the VOLUMENAMES parameter to contain all of the export data (unless PREVIEW=YES). To determine how many volumes you may need, you can run the command specifying PREVIEW=YES.

EXPORT POLICY

VOLumenames=*volname*

Specifies one or more sequential access volumes that are to be used to contain exported files. This parameter is optional. The items in the list are separated by commas, with no intervening spaces. Any volume name used by ADSM cannot contain any embedded blanks or equal signs. A volume name must be specified if SCRATCH=NO and PREVIEW=NO. If you do not specify a volume and SCRATCH=YES, scratch volumes are used. Use the following naming conventions when specifying volumes for each device type.

CARTRIDGE

Specify 1–6 alphanumeric characters, for example DSMT01.

FILE

Specify any supported file name, for example /exndata/mt1.

Examples

Task

Export the policy information related to all domains whose names begin with a D to a device class defined for tape (assuming a device class of CLASS4) and allowing for scratch volumes to be mounted.

The number and type of objects exported are reported to the system console and in the activity log.

Command: EXPort Policy D* DEVclass=class4

Output:

```
ANR0609I EXPORT POLICY started as process 22.  
ANS5104I Process number 22 started.
```

Related Commands

Table 66 provides references to related commands.

Table 66 (Page 1 of 2). Commands Related to EXPORT POLICY

Command	Description	Reference
CANCEL PROCESS	Cancels a background server process.	Page 64
EXPORT ADMIN	Copies administrative information to external media.	Page 195
EXPORT NODE	Copies client node information to external media.	Page 199
EXPORT SERVER	Copies all or part of the server to external media.	Page 208
IMPORT POLICY	Restores policy information from external media.	Page 235

Table 66 (Page 2 of 2). Commands Related to EXPORT POLICY

Command	Description	Reference
QUERY ACTLOG	Searches for and displays messages from the server activity log.	Page 252
QUERY PROCESS	Displays information about one or all background server processes.	Page 332

EXPORT SERVER (Copy Server Information to Sequential Media)

Use the EXPORT SERVER command to export all or part of the server control information and client file data (if specified) from the ADSM server to sequential access media.

The EXPORT SERVER command exports server control information and client node data such as:

- Policy domain definitions
- Policy set definitions
- Management class and copy group definitions
- Schedule definitions and node associations
- Administrator definitions
- Client node definitions

Optionally, the EXPORT SERVER command exports node information such as:

- File space definitions
- Data from files migrated from HSM client nodes
- Archive or backup file copy data
- Access authorization information pertaining to the file spaces reported

In essence, the EXPORT SERVER command allows all global server control information (such as policy data and administrative information) to be copied to another server. In order to backup ADSM data and then restore it to a new server, the device used for export must be supported on both platforms.

ADSM client data can be moved between ADSM servers via export and import processing, if the same removable media type is supported on both platforms.

If a primary file cannot be exported because it is damaged or because it resides on a volume that is UNAVAILABLE or OFFLINE, ADSM attempts to obtain a copy of the file from a copy storage pool.

This command is useful for permitting later retrieval of the server information for copying the information to a new server.

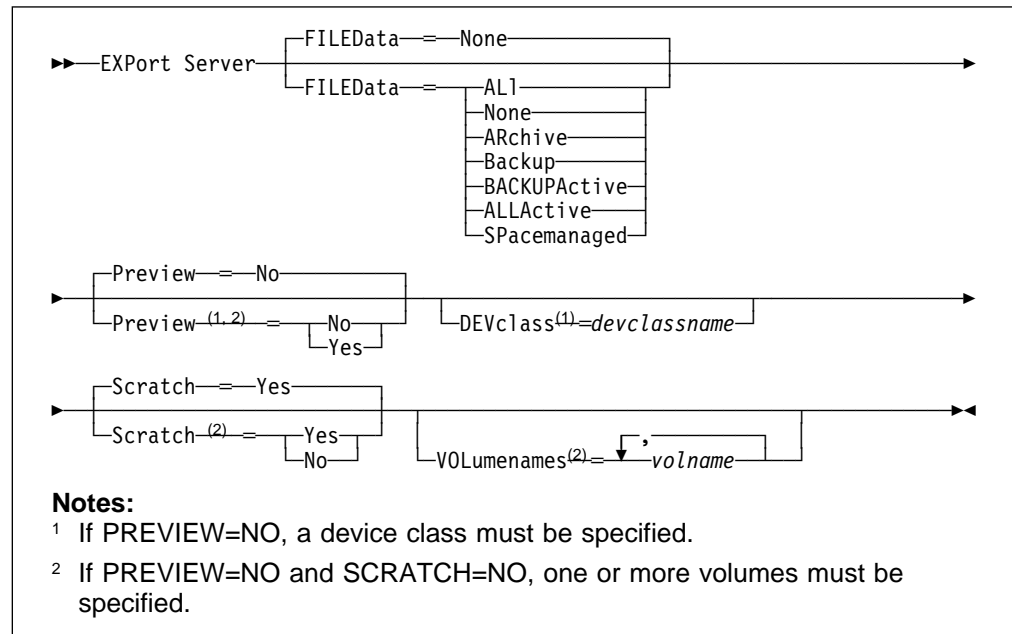
You can use the QUERY ACTLOG command to view the actual status information indicating the size and the success or failure of the export operation. You can also view this information from the server console.

This command generates a background process that can be canceled with the CANCEL PROCESS command. If an EXPORT SERVER background process is canceled, the volumes holding the exported data are incomplete and should not be used for importing data. To display information on background processes, use the QUERY PROCESS command.

Privilege Class

To issue this command, you must have system privilege.

Syntax



Parameters

FILEData=*filedatatype*

Specifies the type of files that should be exported for all nodes defined to the server. This parameter is optional. The default value is NONE.

The device class used by the file data is determined by the device class for the storage pool. If it is the same device class specified in this command, then two drives are needed to export server information. The mount limit value for the device class needs to be set to at least 2.

The following descriptions mention *active* and *inactive* file copies. An active file copy is the most recent backup copy for a file that still exists on the client workstation. All other file copies are called inactive copies. Possible values are:

ALI

Specifies that ADSM exports all backup versions of files, all archived files, and all files migrated from a user's local file system.

None

Specifies that no files are to be exported; only server definitions are to be exported.

ARchive

Specifies that only archived files are to be exported.

Backup

Specifies that only backup versions, whether active or inactive, are to be exported.

BACKUPActive

Specifies that only active backup versions are to be exported.

ALLActive

Specifies that ADSM exports all active backup versions of files, ted. all archived files, and all files migrated from a user's local file system.

SPacemanaged

Specifies that only files that have been migrated from a user's local file system are to be exported.

Preview=*value*

Specifies whether you want to preview the results of the export operation, without actually exporting server information. This command is used to determine how many bytes of data is transferred, allowing you to determine how many sequential access media will be required to hold the data. This parameter is optional. The default value is NO. Possible values are:

Yes

Specifies that you want to preview the results of the export operation, without actually exporting server information. Information on the number and types of objects exported, together with the number of bytes transferred are reported to the server console and the activity log. If you specify this value, you do not need to specify device or volume parameters.

No

Specifies that the server information is to be exported. If you specify this value, you must specify a value for the DEVCLASS parameter.

DEVclass=*devclassname*

Specifies the device class to which export data is to be written. This parameter is required if PREVIEW=NO. You cannot specify the DISK device class.

If the mount limit for the device class specified has already been reached (all drives are busy) when the export runs, ADSM automatically cancels lower priority operations, such as reclamation, to make a mount point available for the export.

Scratch=*scratchvalue*

Specifies whether scratch volumes can be used for export. The default is SCRATCH=YES. Possible values are:

Yes

Specifies that scratch volumes are to be used for export. If you include a list of volumes on the VOLUMENAMES parameter, then scratch volumes are used only if there is not enough space to store the exported files on the volumes specified. If the device type associated with the specified device class is FILE, file names for scratch volumes are generated based on a time stamp.

No

Specifies that scratch volumes cannot be used for export. You must include a list of volumes on the VOLUMENAMES parameter to contain all of the export data, unless PREVIEW=YES. To determine how many volumes you may need, you can run the command specifying PREVIEW=YES.

VOLumenames=*volname*

Specifies one or more sequential access volumes that are to be used to contain exported files. This parameter is optional. The items in the list are

separated by commas, with no intervening spaces. Any volume name used by ADSM cannot contain any embedded blanks or equal signs. A volume name must be specified if SCRATCH=NO and PREVIEW=NO. If you do not specify a volume and SCRATCH=YES, scratch volumes are used. Use the following naming conventions when specifying volumes for each device type.

CARTRIDGE

Specify 1–6 alphanumeric characters, for example DSMT01.

FILE

Specify any supported file name, for example /exndata/mt1.

Examples

Task

Before exporting the server, determine how many objects will be exported, and the total size in bytes required to export the entire server and all archived files.

The number and type of objects exported are reported to the server console and in the activity log.

Command: EXPort Server FILEData=ARchive Preview=Yes

Output:

```
ANR0609I EXPORT SERVER started as process 2.  
ANS5104I Process number 2 started.
```

Related Commands

Table 67 provides references to related commands.

Table 67. Commands Related to EXPORT SERVER

Command	Description	Reference
CANCEL PROCESS	Cancels a background server process.	Page 64
EXPORT ADMIN	Copies administrative information to external media.	Page 195
EXPORT NODE	Copies client node information to external media.	Page 199
EXPORT POLICY	Copies policy information to external media.	Page 204
IMPORT SERVER	Restores all or part of the server from external media.	Page 239
QUERY ACTLOG	Searches for and displays messages from the server activity log.	Page 252
QUERY PROCESS	Displays information about one or all background server processes.	Page 332

EXTEND DB (Increase the Assigned Capacity of the Database)

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. (When you allocate additional volumes for the database, you must add 1MB to the total number of volumes allocated. ADSM uses this additional 1MB for recovery processing.)

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

You can use this command after you define a database volume to increase the assigned capacity of the database. To increase the capacity of the database, you must extend the database in 4MB increments. If you do not specify the extension in 4MB increments, ADSM rounds the number to the next 4MB partition. For example, if you extend the database by 2MB, ADSM extends the capacity of the database by 4MB.

The EXTEND DB command can be initiated while users are accessing the server. Use the QUERY DB command to determine how much you can increase the assigned capacity of the database.

Privilege Class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
▶▶—EXTend DB—nnn————▶▶
```

Parameters

nnn

Specifies the number of megabytes to be added to the assigned capacity of the database. This parameter is required. The value is in increments of 4MB.

Examples

Task

Extend the assigned capacity of the database by 52MB.

Command: EXTend DB 52

Output:

```
ANR2248I Database assigned capacity has been extended.
```

Related Commands

Table 68 provides references to related commands.

Table 68. Commands Related to EXTEND DB

Command	Description	Reference
CANCEL PROCESS	Cancels a background server process.	Page 64
DEFINE DBCOPY	Creates a volume copy of a database volume.	Page 95
DEFINE DBVOLUME	Assigns a volume to be used for database storage.	Page 97
QUERY DB	Displays allocation information about the database.	Page 274
REDUCE DB	Attempts to free up a database volume by reducing usable storage within the volumes defined.	Page 366

EXTEND LOG (Increase the Assigned Capacity of the Recovery 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. (When you allocate additional volumes for the database, you must add 1MB to the total number of volumes allocated. ADSM uses this additional 1MB for recovery processing.)

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

You can use this command after you define a recovery log volume, via the DEFINE LOGVOLUME command, to increase the assigned capacity of the recovery log. To increase the capacity of the recovery log, you must extend the recovery log in 4MB increments. If you do not specify the extension in 4MB increments, ADSM rounds the number to the next 4MB partition. For example, if you extend the recovery log by 1MB, ADSM extends the recovery log by 4MB.

The EXTEND LOG command can be initiated while users are accessing the server. Use the QUERY LOG command to determine how much you can increase the assigned capacity of the recovery log.

This command can generate a background process that can be canceled with the CANCEL PROCESS command. If an EXTEND LOG background process is canceled, the log may be partially extended. To display information on background processes, use the QUERY PROCESS command.

Privilege Class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
▶▶—EXTend LOG—nnn—————▶▶
```

Parameters

nnn

Specifies the number of megabytes to be added to the assigned capacity of the recovery log. This parameter is required. The value is in increments of 4MB.

Examples

Task

Extend the assigned capacity of the recovery log by 20MB.

Command: EXTend LOG 20

EXTEND LOG

Output:

```
ANR2269I Recovery log extension process initiated  
(process id 1).  
ANS5104I Process number 1 started.
```

Related Commands

Table 69 provides references to related commands.

Table 69. Commands Related to EXTEND LOG

Command	Description	Reference
CANCEL PROCESS	Cancels a background server process.	Page 64
DEFINE LOGCOPY	Creates a volume copy of a recovery log volume.	Page 110
DEFINE LOGVOLUME	Defines a disk volume to the recovery log.	Page 112
QUERY LOG	Displays allocation information about the recovery log.	Page 307
QUERY PROCESS	Displays information about one or all background server processes.	Page 332
REDUCE LOG	Attempts to free up a log volume by reducing usable storage within the volumes defined.	Page 368

GRANT AUTHORITY (Add Administrator Authority)

Use the GRANT AUTHORITY command to grant an administrator one or more administrative privilege classes.

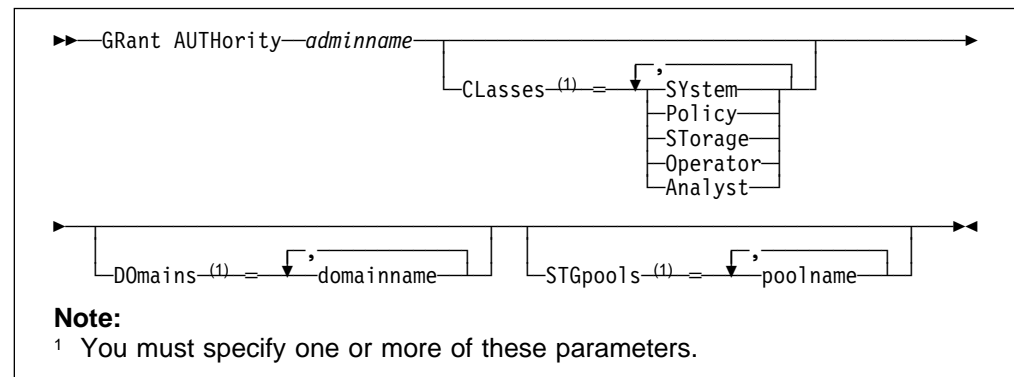
You can use this command to grant administrative privilege classes to an administrator that has not been granted any administrative privilege, or to add additional privilege classes to an administrator.

You cannot grant restricted privilege to an unrestricted policy or storage administrator. To do this, you must use the REVOKE AUTHORITY command to remove the administrator's unrestricted privilege, then use this command to grant restricted privilege to the administrator.

Privilege Class

To issue this command you must have system privilege.

Syntax



Parameters

adminname

Specifies the name of the administrator to receive administrative authority. This parameter is required.

CLASSES=classes

Specifies one or more privilege classes to be granted to an administrator. You can specify more than one privilege class by delimiting each with a comma. Possible classes are:

SYstem

Specifies that you want to grant system privilege to an administrator. Do not specify additional privilege classes or the DOMAINS or STGPOLS parameters when granting system privilege to an administrator. A system administrator has the highest level of authority in ADSM. A system administrator can issue any administrative command and has authority to manage all policy domains and all storage pools. Only a system administrator can grant authority to other administrators.

Policy

Specifies that you want to grant policy privilege to an administrator. If the DOMAINS parameter is not specified, unrestricted policy privilege is granted. An unrestricted policy administrator can issue commands that

GRANT AUTHORITY

affect all existing policy domains as well as any policy domains that are defined in the future. An unrestricted policy administrator cannot define, delete, or copy policy domains. A restricted policy administrator can be upgraded to an unrestricted policy administrator by using the GRANT AUTHORITY command with CLASSES=POLICY and no DOMAINS parameter.

Storage

Specifies that you want to grant storage privilege to an administrator. If the STGPools parameter is not specified, unrestricted storage privilege is granted. An unrestricted storage administrator can issue all commands that allocate and control storage resources for the server. An unrestricted storage administrator can issue commands that affect all existing storage pools as well as any storage pools that are defined in the future. An unrestricted storage administrator cannot define or delete storage pools. A restricted storage administrator can be upgraded to an unrestricted storage administrator by using the GRANT AUTHORITY command with CLASSES=STORAGE and no STGPools parameter.

Operator

Specifies that you want to grant operator privilege to an administrator. An administrator with operator privilege can issue commands that control the immediate operation of the server and the availability of storage media.

Analyst

Specifies that you want to grant analyst privilege to an administrator. An administrator with analyst privilege can issue commands that reset the counters which track server statistics.

DOMains=*domainname*

Specifies that you want to grant restricted policy privilege to an administrator. If the DOMAINS parameter is specified, then CLASSES=POLICY is optional.

Restricted policy privilege permits an administrator to issue a subset of the policy commands for the domains to which the administrator is authorized. You can use this parameter to grant additional policy domain authority to a restricted policy administrator. This parameter is optional. You can specify more than one policy domain by delimiting each policy domain name with a comma.

You can use a pattern-matching expression to specify a name. Authority for all matching policy domains is granted.

STGpools=*poolname*

Specifies that you want to grant restricted storage privilege to an administrator. If the STGPools parameter is specified, then CLASSES=STORAGE is optional.

Restricted storage privilege permits you to issue a subset of the storage commands for the storage pools to which the administrator is authorized. You can use this parameter to grant additional storage pool authority to a restricted storage administrator. This parameter is optional. You can specify more than one storage pool by delimiting each storage pool name with a comma.

You can use a pattern-matching expression to specify a name. Authority for all matching storage pools is granted.

Examples

Task 1

Grant system privilege to administrator BILL.

Command: GRant AUTHority bill CLasses=SYstem

Output:

```
ANR2076I System privilege granted to administrator BILL.
```

Task 2

Specify additional policy domains that can be managed by the restricted policy administrator CLAUDIA.

Command: GRant AUTHority claudia DOmains=domain3,domain4

Output:

```
ANR2078I Restricted policy privilege granted to
administrator CLAUDIA - policy domain DOMAIN3.
ANR2078I Restricted policy privilege granted to
administrator CLAUDIA - policy domain DOMAIN4.
```

Task 3

Provide administrator TOM with unrestricted storage privilege and restricted policy privilege for the domains whose names start with AIX.

Command: GRant AUTHority tom CLasses=STorage DOmains=AIX*

Output:

```
ANR2078I Restricted policy privilege granted to
administrator TOM - policy domain AIX3.
ANR2078I Restricted policy privilege granted to
administrator TOM - policy domain AIX4.
ANR2079I Unrestricted storage privilege granted
to administrator TOM.
```

Related Commands

Table 70 provides references to related commands.

Table 70. Commands Related to GRANT AUTHORITY

Command	Description	Reference
QUERY ADMIN	Displays information about one or more ADSM administrators.	Page 256
REVOKE AUTHORITY	Revokes one or more privilege classes or restricts access to policy domains and storage pools.	Page 401

HALT (Shut down the Server)

Use the HALT command to shut down the server. The HALT command forces an abrupt shutdown and results in the canceling of all administrative and client node sessions, even if they are not completed.

Any in-progress transactions interrupted by this command are rolled back when you restart the server. If the QUIESCE option is specified, transaction activity is quiesced when current transactions have completed, and then the server is shutdown.

The QUIESCE option is recommended ONLY if you plan to perform a stand-alone database dump by using the DSMSERV DUMPDB command immediately following server shutdown. For additional information about the DSMSERV DUMPDB command and related commands, see Appendix B, “Database Salvage Commands” on page 529.

It is recommended to use the HALT command only after all administrative and client node sessions are completed or canceled. In order to shut down the server without severely impacting administrative and client node sessions, you can:

1. Use the DISABLE command to prevent new client node sessions from starting.
2. Use the QUERY SESSIONS command to identify any existing administrative and client node sessions.
3. Notify any existing administrative and client node sessions that you plan to shut down the server (you must do this outside of ADSM).
4. Use the CANCEL SESSIONS command to cancel any existing administrative or client node sessions.
5. Issue the HALT command to shut down the server and stop any administrative and client node sessions.

To start the server after issuing the HALT command, see the *ADSM Administrator's Guide*.

Privilege Class

To issue this command, you must have system privilege or operator privilege.

Syntax

```
▶▶ HALT [Quiesce] ▶▶
```

Parameters

Quiesce

Specifies that all database transaction activity should be gracefully stopped when current transactions complete. This option should only be used when you are planning to perform a stand-alone database dump using the DSMSERV DUMPDB command to dump database contents to sequential media, such as tape.

The default results with any in-progress transactions interrupted by this command are rolled back when you restart the server.

HALT

When using the QUIESCE option on the HALT command, the server attempts to stop all database update transactions against the database and record database pages to stable media so that a consistent image of the database can be dumped with the DSMSERV DUMPDB command. This attempt to stop update activity may be retried a number of times, and may even fail if there are many transactions in progress on the server. If this failure occurs the HALT command stops processing and the server continues to run. If this occurs, please cancel all client sessions, cancel all background server processes, and re-issue the HALT QUIESCE command. The HALT command without the QUIESCE option will always cause the server to halt successfully.

Examples

Task 1

From the server console, shut down the server abruptly. All user activity stops immediately and no new activity is allowed to start.

Command: HALT

Output:

```
ANR0991I ADSM server shut down complete.
```

Task 2

From an administrative client, shut down the server abruptly. All user activity stops immediately and no new activity is allowed to start.

Command: HALT

Output:

```
ANR2234W This command will halt the server; if the command
is issued from a remote client, it may not be possible to
restart the server from the remote location.
Do you wish to proceed? (Y/N)
```

Type Y and press the ENTER key. A message similar to the following is displayed:

```
ANS5103I Highest return code was 50.
```

Task 3

When planning to perform the stand-alone database dump using the DSMSERV DUMPDB command, shut down the ADSM server with the QUIESCE option.

Command: HALT Quiesce

Output:

```

ANR2106I HALT: Quiescing database update activity.
ANR2107I HALT: Database update activity is now quiesced.
ANR0991I ADSM server shutdown complete.
ANR5963I ADSM server termination complete.
    
```

The stand-alone dump can now be performed.

Related Commands

Table 71 provides references to related commands.

Table 71. Commands Related to HALT

Command	Description	Reference
CANCEL SESSION	Cancels one or all active sessions with the server.	Page 66
DISABLE	Prevents new clients from accessing ADSM but permits existing sessions to continue.	Page 188
ENABLE	Resumes server activity following the DISABLE command.	Page 192
QUERY SESSION	Displays information about all active administrator and client sessions with ADSM.	Page 339

HELP (Get Help on Commands and Error Messages)

Use the HELP command to display information about administrative commands and error messages from the server and administrative command line client. For the best usability, issue the HELP command from an administrative command line client.

You can access the help facility by entering one of the following command strings:

- Help
- Help *help_menu_number*
- Help *commandname*
- Help *commandname subcommandname*
- Help *error_msg_number*

The HELP command (with no operands) displays a menu with a list of selections that provide information on using the command line interface and a list of the ADSM administrative commands.

The Help *help_menu_number* command displays the help text for a specific menu selection from the help menu.

The Help *commandname* command displays the help text specific to that command. If this command has subcommand names, you will also receive a list of the subcommand names. You can then select specific help related to the subcommand name.


The Help *commandname subcommandname* command displays the help text specific to the command and its associated sub-commands.

The Help *error_msg_number* command displays the help text specific to the error message. You can get help information on server messages (prefixed by ANR) and client messages (prefixed by ANS).

Privilege Class

Any administrator can issue this command.

Syntax



Parameters

help_menu_number

Specifies the number of your selection from the help menu. This parameter is optional.

commandname

Specifies the name of the administrative command you want to display. This parameter is optional.

subcommandname

Specifies the subcommand name associated with the name of administrative command that you want to display. This parameter is optional.

error_msg_number

Specifies the number of the error message for which you want to display information. This parameter is optional. You can get help information on server messages (prefixed by ANR) and client messages (prefixed by ANS). You do not have to include the prefix and severity code when specifying an error message number.

Examples

Task 1

Display the help menu for the command-line interface.

Command:Help

Output:

```
The following help sections are available:
 1 - Using Syntax Diagrams
 2 - Entering Administrative Commands
 3 - Using the Server Console Session
 4 - Using the Administrative Client
 5 - Using the Administrative Client Options
 6 - Using Continuation Characters
 7 - Using Console Mode
 8 - Using Batch Mode
 9 - Using Interactive Mode
10 - Redirecting Command Output to a File
11 - Using Macros to Issue Commands
12 - Using Commands Based on Privilege Classes
13 - Commands Requiring System Privilege
14 - Commands Requiring Policy Privilege
15 - Commands Requiring Storage Privilege
16 - Commands Requiring Operator Privilege
17 - Commands Requiring Analyst Privilege
18 - Commands Any Administrator Can Issue
```

Task 2

Display help information using the help menu in Example 1.

Command:Help 18

Output:

Information on commands that can be issued by any administrator is displayed.

Task 3

Display help information on the REMOVE commands.

Command: Help REMove

Output:

```
REMOVE
-----
Use the REMOVE commands to remove an object from ADSM.
The following is a list of REMOVE commands for ADSM:
    "REMOVE ADMIN (Delete an Administrator)"
    "REMOVE NODE (Delete a Client Node)"
```

Task 4

Display help information on the DEFINE COPYGROUP command.

Command: Help DEFine C0pygroup

Output:

Information on the DEFINE COPYGROUP command, on page 83 is displayed.

Task 5

Display help information on the error message ANR5300I.

Command: Help 5300

Output:

```
-----
ANR5300I  VERIFYING VOLUME LABEL volume ID ON DEVICE real address.
EXPLANATION: The server is reverifying the volume label of a mounted
read/write volume just prior to demounting the volume.
SYSTEM ACTION: The server performs the reverification of the volume label as
an integrity check of the volume.
USER RESPONSE: None.
```

Related Commands

None.

IMPORT ADMIN (Import Administrator Information)

Use the IMPORT ADMIN command to import administrator and authority definitions for one or more administrators from export media to the ADSM server.

The IMPORT ADMIN command imports administrator information such as:

- Administrator name, password, and contact information
- Administrative privilege classes the administrator has been granted

The IMPORT ADMIN command allows administrative information to be copied to another server. In order to copy ADSM data and then restore it to a new server, the sequential access device must be supported on both platforms.

ADSM client data can be moved between ADSM servers via export and import processing, if the same removable media type is supported on both platforms.

This command is useful for permitting later retrieval of server information for copying the information to a new server.

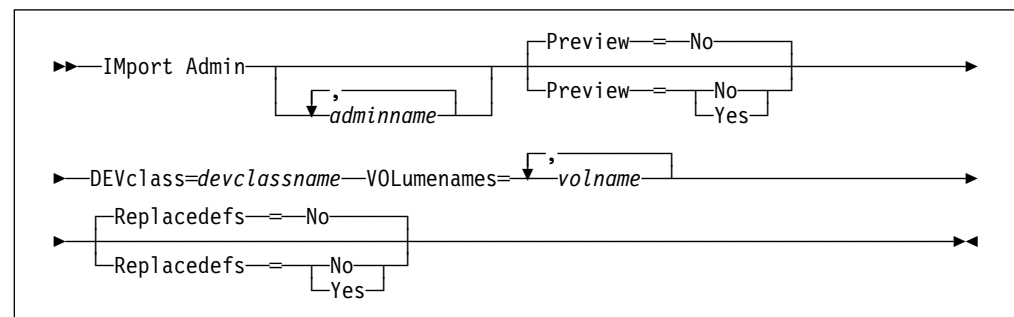
You can use the QUERY ACTLOG command to view the actual status information indicating the size and the success or failure of the import operation. You can also view this information from the server console.

This command generates a background process that can be canceled with the CANCEL PROCESS command. To display information on background processes, use the QUERY PROCESS command.

Privilege Class

To issue this command, you must have system privilege.

Syntax



Parameters

adminname

Specifies the name of an administrator for which you want import administrator information. This parameter is optional.

You can specify a list of administrators for which administrator information is to be imported. The items in the list are separated by commas, with no intervening spaces. You can use pattern-matching expressions to specify names. All matching administrators are included in the list. If you do not specify a value for this parameter, information related to all administrators found on the import media is imported.

Preview=*value*

Specifies whether you want to preview the results of the import operation, without actually importing administrator information. For the import commands, the PREVIEW=YES option may require that export tapes be mounted to generate the preview statistics. This parameter is optional. The default value is NO. Possible values are:

Yes

Specifies that you want to preview the results of the import operation, without actually importing administrator information. Information on the number and types of objects imported, together with the number of bytes transferred are reported to the server console and the activity log.

No

Specifies that the information is to be imported.

DEVclass=*devclassname*

Specifies the device class from which import data is to be read. This parameter is required. You cannot specify the DISK device class.

If the mount limit for the device class specified has already been reached (all drives are busy) when the import runs, ADSM automatically cancels lower priority operations, such as reclamation, to make a mount point available for the import.

Replacedefs=*value*

Specifies whether node definitions which exist on the server should be replaced by imported objects having the same name. Existing file spaces are not replaced. New file spaces are created when identical file space names are encountered. The default value is NO. Possible values are:

Yes

Specifies that definitions are to be replaced by the imported objects.

Warning: When specifying REPLACEDFS=YES, be sure that you want to import data.

No

Specifies that definitions are not to be replaced by imported objects.

VOLumenames=*volname*

Specifies the name of a volume from which you want to import information. You can specify a list of volumes for which volume information is to be imported. The list of volumes needs to be in the same sequential order as they were exported in. The items in the list are separated by commas, with no intervening spaces. Any volume name used by ADSM cannot contain any embedded blanks or equal signs. Use the following naming conventions when specifying volumes for each device type.

CARTRIDGE

Specify 1–6 alphanumeric characters, for example DSMT01.

FILE

Specify any supported file name, for example /exndata/mt1.

Examples

Task

Import the administrator information for all defined administrators from a cartridge tape (assuming a device class of 3480CART).

The number and type of objects imported are reported to the system console and in the activity log.

Command: `IMport Admin DEVclass=3480CART`

Output:

```
ANR0609I IMPORT ADMIN started as process 19.
ANS5104I Process number 19 started.
```

Related Commands

Table 72 provides references to related commands.

Table 72. Commands Related to IMPORT ADMIN

Command	Description	Reference
CANCEL PROCESS	Cancels a background server process.	Page 64
EXPORT ADMIN	Copies administrative information to external media.	Page 195
IMPORT NODE	Restores client node information from external media.	Page 230
IMPORT POLICY	Restores policy information from external media.	Page 235
IMPORT SERVER	Restores all or part of the server from external media.	Page 239
QUERY ACTLOG	Searches for and displays messages from the server activity log.	Page 252
QUERY PROCESS	Displays information about one or all background server processes.	Page 332

IMPORT NODE (Import Client Node Information)

Use the IMPORT NODE command to import client node definitions from the ADSM server to sequential media. Client node definitions include information such as:

- User ID, password, and contact information
- Name of the policy domain to which the client belongs
- File compression status
- Whether the user has the authority to delete backed up or archive files from data storage
- Whether the client node ID was locked at the time of export

Optionally, IMPORT NODE command imports node data information such as:

- File space definitions
- Data from files migrated from HSM client nodes
- Archive or backup file copy data
- Access authorization information pertaining to the file spaces reported

The IMPORT NODE command allows node information to be copied to another server. In order to copy ADSM data and then restore it to a new server, the sequential access device must be supported on both platforms.

This command is useful for permitting later retrieval of server information for copying the information to a new server.

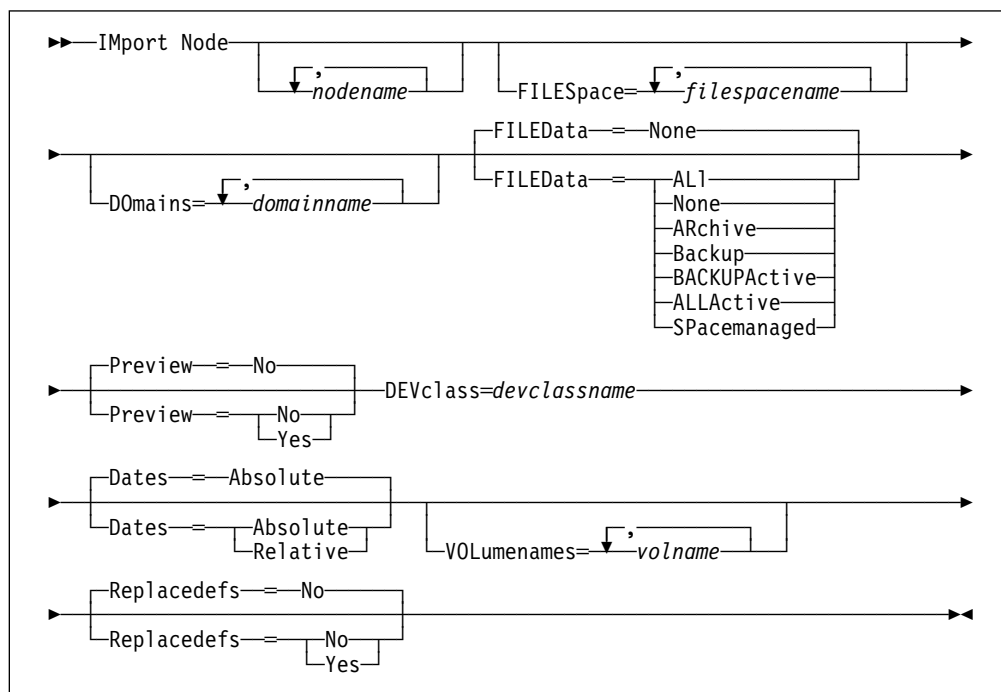
You can use the QUERY ACTLOG command to view the actual status information indicating the size and the success or failure of the import operation. You can also view this information from the server console.

This command generates a background process that can be canceled with the CANCEL PROCESS command. If an IMPORT NODE background process is canceled, some of the data may have already been imported. If so, you may want to delete the partially imported file spaces. To display information on background processes, use the QUERY PROCESS command.

Privilege Class

To issue this command, you must have system privilege.

Syntax



Parameters

nodename

Specifies the name of a client node for which you want import client node information. This parameter is optional. You can specify a list of node names for which node information is to be imported. The items in the list are separated by commas, with no intervening spaces. You can use pattern-matching expressions to specify names. All matching nodes are included in the list. If you do not specify a value for this parameter, information related to all nodes is imported.

FILESpace=filespace

Specifies the name of a file space for which you want import file space information. This parameter is optional. You can specify a list of file space names for which data information is to be imported. The items in the list are separated by commas, with no intervening spaces.

You can use a pattern-matching expression to specify a name. All matching names are included in the list. If you do not specify a value for this parameter, information related to all file space names is imported.

Note: Existing file spaces are not replaced. New file spaces are created when identical names are encountered.

DOMAINS=domainname

Specifies the names of the policy domains to which nodes must be assigned in order for node information to be imported. This parameter is optional. You can specify a list of policy domain for which node information is to be imported. The items in the list are separated by commas, with no intervening spaces. You can use a pattern-matching expression to specify a name. All matching domains are included in the list. If you do not specify a value for this parameter, node information related to all domains is imported.

FILEData=type

Specifies the type of files that should be imported for all nodes specified and found on the export volumes. This parameter is optional. The default value is NONE.

The device class used by the file data is determined by the device class for the storage pool. If it is the same device class specified in this command, then two drives are needed to import node information. The mount limit value for the device class needs to be set to at least 2.

The following descriptions mention *active* and *inactive* file copies. An active file copy is the most recent backup copy for a file that still exists on the client workstation. All other file copies are called inactive copies. Possible values are:

ALI

Specifies that ADSM imports all backup versions of files, all archived, and all files migrated from a user's local file system as qualified by the NODE and DOMAINS parameters.

None

Specifies that no files are to be imported; only node definitions are to be imported.

ARChive

Specifies that only archived files are to be imported.

Backup

Specifies that only backup versions, whether active or inactive, are to be imported.

BACKUPActive

Specifies that only active backup versions are to be imported.

ALLActive

Specifies that ADSM imports all active backup versions of files, all archived, and all files migrated from a user's local file system.

SPacemanaged

Specifies that only files that have been migrated from a user's local file system are to be imported.

Preview=value

Specifies whether you want to preview the results of the import operation, without actually importing information. For the import commands, the PREVIEW=YES option may require that all export volumes be mounted to generate the preview statistics. This parameter is optional. The default value is NO. Possible values are:

Yes

Specifies that you want to preview the results of the import operation, without actually importing files. Information on the number and types of objects imported, together with the number of bytes transferred are reported to the server console and the activity log.

No

Specifies that the node information is to be imported.

DEVclass=devclassname

Specifies the device class to which import data is to be read. You cannot specify the DISK device class. This parameter is required.

If the mount limit for the device class specified has already been reached (all drives are busy) when the import runs, ADSM automatically cancels lower priority operations, such as reclamation, to make a mount point available for the import.

Dates=value

Specifies whether the recorded backup or archive dates for client node file copies are set to the values specified when the files were exported (ABSOLUTE), or are adjusted relative to the date of import (RELATIVE). The default value is ABSOLUTE.

If the export media has been idle for some time after export (sitting on a shelf for six months, for example), the original backup or archive dates may be old enough to trigger immediate roll-off of file copies when the data is imported into a server. The Relative specification for this value will adjust for time elapsed since export so that the file copies are not immediately rolled off.

For example, assume that an export tape contains an archive file copy that was archived 5 days prior to the export operation. If the export media is saved for 6 months and then imported, the archive file will appear to have been inserted six months and 5 days ago by default (DATES=ABSOLUTE) and may be subject to immediate roll-off depending upon the retention value that is specified in the file's management class. Specifying DATES=RELATIVE results in resetting the archive date for the file to 5 days ago during import, even though the actual file was archived six months and 5 days ago. The DATES parameter then adjusts file backup and archive dates for the time that has elapsed since the export operation occurred.

Possible values are:

ABSOLUTE

Specifies the date for file copies are set to the values specified when the files were exported.

RELATIVE

Specifies the date for file copies are adjusted to the date of import.

VOLumentnames=volname

Specifies the name of a volume from which you want to import volume information. This parameter is required. You can specify a list of volumes from which volume information is to be imported. The items in the list are separated with commas, with no intervening spaces. Use the following naming conventions when specifying volumes for each device type.

CARTRIDGE

Specify 1–6 alphanumeric characters, for example DSMT01.

FILE

Specify any supported file name, for example /exndata/mt1.

Replacedefs=value

Specifies whether node definition which exist on the server should be replaced by imported objects having the same name. Existing file spaces are not replaced. New file spaces are created when identical file space names are encountered. The default value is NO. Possible values are:

IMPORT NODE

Yes

Specifies that objects are to be replaced by the imported objects.

Warning: When specifying REPLACEDFS=YES, be sure that you want to import data.

No

Specifies that objects are not to be replaced by imported objects.

Examples

Task

Import the node definitions and active backup data for all nodes defined in the STANDARD policy domain from tape (assuming a device class of DEPT).

The number and type of objects imported are reported to the server console and in the activity log.

Command: `Import Node D0mains=standard DEVclass=DEPT
FILEData=BACKUPActive`

Output:

```
ANR0609I IMPORT NODE started as process 32.  
ANS5104I Process number 32 started.
```

Related Commands

Table 73 provides references to related commands.

Table 73. Commands Related to IMPORT NODE

Command	Description	Reference
CANCEL PROCESS	Cancels a background server process.	Page 64
EXPORT NODE	Copies client node information to external media.	Page 199
IMPORT ADMIN	Restores administrative information from external media.	Page 227
IMPORT POLICY	Restores policy information from external media.	Page 235
IMPORT SERVER	Restores all or part of the server from external media.	Page 239
QUERY ACTLOG	Searches for and displays messages from the server activity log.	Page 252
QUERY PROCESS	Displays information about one or all background server processes.	Page 332

IMPORT POLICY (Import Policy Information)

Use the IMPORT POLICY command to import policy information for one or more policy domains from export media to the ADSM server.

The IMPORT POLICY command imports policy information and client node data such as:

- Policy domain definitions
- Policy set definitions, including the active policy set
- Management class definitions, including the default management class
- Backup copy group and archive copy group definitions
- Schedule definitions and client node associations

The IMPORT POLICY command allows policy information to be copied to another server. In order to backup ADSM data and then restore it to a new server, the device must be supported on both platforms.

ADSM client data can be moved between ADSM servers via export and import processing, if the same removable media type is supported on both platforms.

This command is useful for permitting later retrieval of the server information for copying the information to a new server.

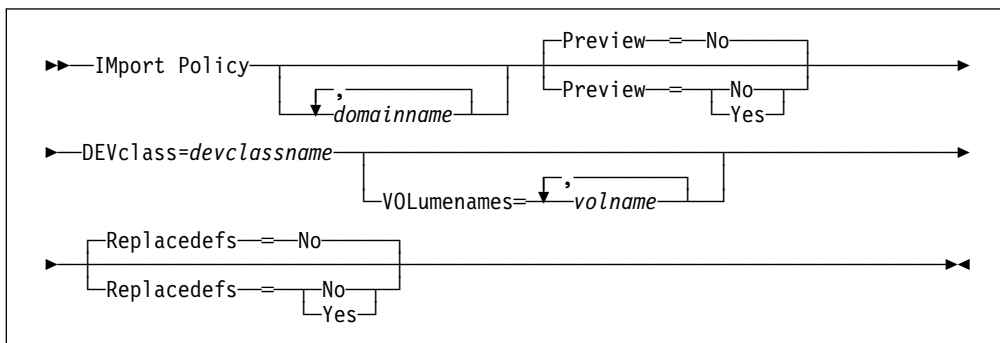
You can use the QUERY ACTLOG command to view the actual status information indicating the size and the success or failure of the import operation. You can also view this information from the server console.

This command generates a background process that can be canceled with the CANCEL PROCESS command. To display information on background processes, use the QUERY PROCESS command.

Privilege Class

To issue this command, you must have system privilege.

Syntax



Parameters

domainname

Specifies a list of policy domains for which policy information is to be imported. If you specify more than one policy domain, separate the names with commas, with no intervening spaces. You can use pattern-matching expressions to specify names. The default (*) is to import the information related to all policy domains defined on the export media.

Preview=*value*

Specifies whether you want to preview the results of the import operation, without actually importing information. For the import commands, the PREVIEW=YES option may require that all export volumes be mounted to generate the preview statistics. This parameter is optional. The default value is NO. Possible values are:

Yes

Specifies that you want to preview the results of the import operation, without actually importing files. Information on the number and types of objects imported, together with the number of bytes transferred are reported to the server console and the activity log.

No

Specifies that the information is to be imported.

DEVclass=*devclassname*

Specifies the device class to which import data is to be read. You cannot specify the DISK device class. This parameter is required.

If the mount limit for the device class specified has already been reached (all drives are busy) when the import runs, ADSM automatically cancels lower priority operations, such as reclamation, to make a mount point available for the import operation.

Replacedefs=*value*

Specifies whether node definitions that exist on the server should be replaced by imported objects having the same name. Existing file spaces are not replaced. New file spaces are created when identical file space names are encountered. The default value is NO. Possible values are:

Yes

Specifies that objects are to be replaced by the imported objects.

Warning: When specifying REPLACEDFS=YES, be sure that you want to import data.

No

Specifies that objects are not to be replaced by imported objects.

VOLumenames=volname

Specifies the name of a volume from which you want to import volume information. This parameter is required. You can specify a list of volumes from which volume information is to be imported. The items in the list are separated by commas, with no intervening spaces. Use the following naming conventions when specifying volumes for each device type.

CARTRIDGE

Specify 1–6 alphanumeric characters, for example DSMT01.

FILE

Specify any supported file name, for example /exndata/mt1.

Examples

Task

Import the policy information related to all domains whose names begin with a D from 3 volumes (EXP001, EXP002, EXP003) on 3480 cartridge tape (assuming a device class of CART).

The number and type of objects imported are reported to the system console and in the activity log.

Command: `Import Policy d* DEVclass=cart VOLumenames=exp001,exp002,exp003`

Output:

```
ANR0609I IMPORT POLICY started as process 21.
ANS5104I Process number 21 started.
```

Related Commands

Table 74 provides references to related commands.

Table 74 (Page 1 of 2). Commands Related to IMPORT POLICY

Command	Description	Reference
CANCEL PROCESS	Cancels a background server process.	Page 64
EXPORT POLICY	Copies policy information to external media.	Page 204
IMPORT ADMIN	Restores administrative information from external media.	Page 227
IMPORT NODE	Restores client node information from external media.	Page 230
IMPORT SERVER	Restores all or part of the server from external media.	Page 239

IMPORT POLICY

Table 74 (Page 2 of 2). Commands Related to IMPORT POLICY

Command	Description	Reference
QUERY ACTLOG	Searches for and displays messages from the server activity log.	Page 252
QUERY PROCESS	Displays information about one or all background server processes.	Page 332

IMPORT SERVER (Import the Server)

Use the IMPORT SERVER command to copy all or part of the server control information and client file data (if specified) from export media to the ADSM server.

The IMPORT SERVER command imports server control information and client node data such as:

- Policy domain definitions
- Policy set definitions
- Management class and copy group definitions
- Schedule definitions and node associations
- Administrator definitions
- Client node definitions

Optionally, IMPORT SERVER command imports server data such as:

- File space definitions
- Data from files migrated from HSM client nodes
- Archive or backup file copy data
- Access authorization information pertaining to the file spaces reported

The IMPORT SERVER command allows all global server control information (such as policy data and administrative information) to be copied to another server. In order to backup ADSM data and then restore it to a new server, the device must be supported on both platforms.

ADSM client data can be moved between ADSM servers via export/import processing, if the same removable media type is supported on both platforms.

This command is useful for permitting later retrieval of the server information for copying the information to a new server.

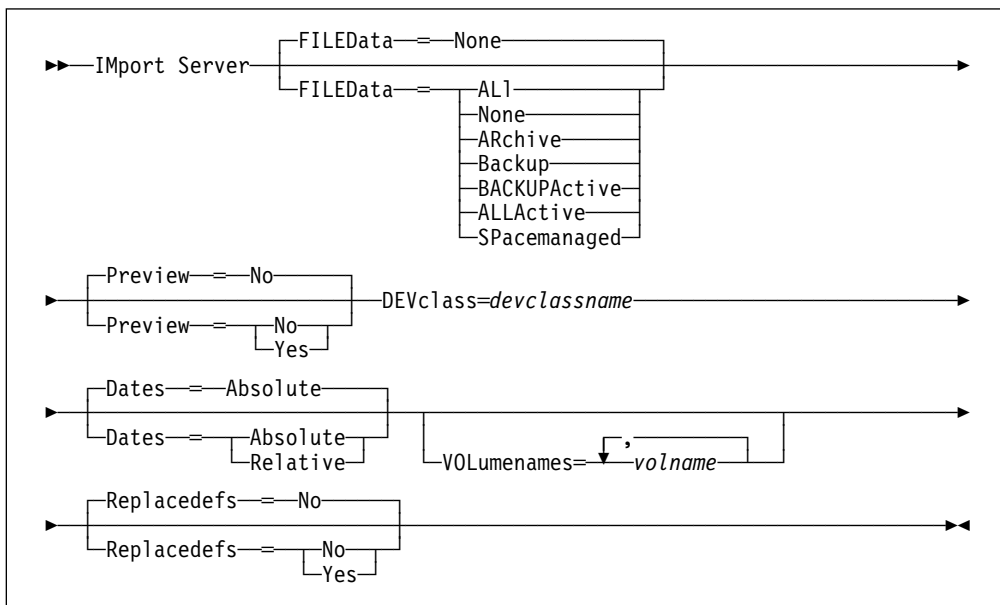
You can use the QUERY ACTLOG command to view the actual status information indicating the size and the success or failure of the import operation. You can also view this information from the server console.

This command generates a background process that can be canceled with the CANCEL PROCESS command. If an IMPORT SERVER background process is canceled, some of the data may have already been imported. If so, you may want to delete the partially imported file spaces. To display information on background processes, use the QUERY PROCESS command.

Privilege Class

To issue this command, you must have system privilege.

Syntax



Parameters

FILEData=*type*

Specifies the type of files that should be imported for all nodes defined to the server. This parameter is optional. The default value is NONE.

The device class used by the file data is determined by the device class for the storage pool. If it is the same device class specified in this command, then two drives are needed to import server information. The mount limit value for the device class needs to be set to at least 2.

The following descriptions mention *active* and *inactive* file copies. An active file copy is the most recent backup copy for a file that still exists on the client workstation. All other file copies are called inactive copies. Possible values are:

ALI

Specifies that ADSM imports all backup versions of files, all archived, and all files migrated from a user's local file system as qualified by the NODE and DOMAINS parameters.

None

Specifies that no files are to be imported; only server definitions are to be imported.

ARchive

Specifies that only archived files are to be imported.

Backup

Specifies that only backup versions, whether active or inactive, are to be imported.

BACKUPActive

Specifies that only active backup versions are to be imported.

ALLActive

Specifies that ADSM imports all active backup versions of files, all archived, and all files migrated from a user's local file system.

SPacemanaged

Specifies that only files that have been migrated from a user's local file system are to be imported.

Preview=*value*

Specifies whether you want to preview the results of the import operation, without actually importing server information. For the import commands, the PREVIEW=YES option may require that all export volumes be mounted to generate the preview statistics. This parameter is optional. The default value is NO. Possible values are:

Yes

Specifies that you want to preview the results of the import operation, without actually importing files. Information on the number and types of objects imported, together with the number of bytes transferred are reported to the server console and the activity log.

No

Specifies that the server information is to be imported.

DEVclass=*devclassname*

Specifies the device class to which import data is to be read. You cannot specify the DISK device class. This parameter is required.

If the mount limit for the device class specified has already been reached (all drives are busy) when the import runs, ADSM automatically cancels lower priority operations, such as reclamation, to make a mount point available for the import operation.

Dates=*value*

Specifies whether the recorded backup or archive dates for client node file copies are set to the values specified when the files were exported (Absolute), or are adjusted relative to the date of import (Relative). The default value is ABSOLUTE.

If the export media has been idle for some time after export (sitting on a shelf for six months, for example), the original backup or archive dates may be old enough to trigger immediate roll-off of file copies when the data is imported into a server. The Relative specification for this value will adjust for time elapsed since export so that the file copies are not immediately rolled off.

For example, assume that an export tape contains an archive file copy that was archived 5 days prior to the export operation. If the export media is saved for 6 months and then imported, the archive file will appear to have been inserted six months and 5 days ago by default (Dates=Absolute) and may be subject to immediate roll-off depending upon the retention value that is specified in the file's management class. Specifying Dates=Relative will result in resetting the archive date for the file to 5 days ago during import, even though the actual file was archived six months and 5 days ago. The DATES parameter then adjusts file backup/archive dates for the time that has elapsed since the export operation occurred.

IMPORT SERVER

Possible values are:

ABSOLUTE

Specifies the date for file copies are set to the values specified when the files were exported.

RELATIVE

Specifies the date for file copies are adjusted to the date of import.

VOLnames=*volname*

Specifies the name of a volume from which you want to import volume information. This parameter is required. You can specify a list of volumes from which volume information is to be imported. The items in the list are separated, with no intervening spaces. Use the following naming conventions when specifying volumes for each device type.

CARTRIDGE

Specify 1–6 alphanumeric characters, for example DSMT01.

FILE

Specify any supported file name, for example /exndata/mt1.

Replacedefs=*value*

Specifies whether node definitions that exist on the server should be replaced by imported objects having the same name. Existing file spaces are not replaced. New file spaces are created when identical file space names are encountered. The default value is NO. Possible values are:

Yes

Specifies that objects are to be replaced by the imported objects.

Warning: When specifying REPLACEDFS=YES, be sure that you want to import data.

No

Specifies that objects are not to be replaced by imported objects.

Examples

Task

Before importing the server control information, determine how many objects are imported and the total number of bytes copied to import the server control information and all archived files from three volumes (exp001,exp002,exp003) on 3480 cartridge tape.

The number and type of objects imported is reported to the server console and in the activity log.

Command: `Import Server FILEData=ARchive DEVclass=cart
VOLumenames=exp001,exp002,exp003 Preview=Yes`

Output:

```
ANR0609I IMPORT SERVER started as process 2.  
ANS5104I Process number 2 started.
```

Related Commands

Table 75 provides references to related commands.

Table 75. Commands Related to IMPORT SERVER

Command	Description	Reference
CANCEL PROCESS	Cancels a background server process.	Page 64
EXPORT SERVER	Copies all or part of the server to external media.	Page 208
IMPORT ADMIN	Restores administrative information from external media.	Page 227
IMPORT NODE	Restores client node information from external media.	Page 230
IMPORT POLICY	Restores policy information from external media.	Page 235
QUERY ACTLOG	Searches for and displays messages from the server activity log.	Page 252
QUERY PROCESS	Displays information about one or all background server processes.	Page 332

LOCK ADMIN (Lock Out an Administrator)

Use the LOCK ADMIN command to prevent an administrator from accessing the server. By using the LOCK ADMIN and UNLOCK ADMIN commands, you can temporarily restrict an administrator from accessing the server. The administrator is locked out until a system administrator uses the UNLOCK ADMIN command to reestablish access for the administrator.

You cannot issue the LOCK ADMIN command against the SERVER_CONSOLE administrative ID.

Privilege Class

To issue this command, you must have system privilege.

Syntax

```
▶▶—LOCK Admin—adminname————▶▶
```

Parameters

adminname
Specifies the name of the administrator to be locked out. This parameter is required.

Examples

Task

Lock out the administrator JOE.

Command: LOCK Admin joe

Output:

```
ANR2072I Administrator JOE locked.
```

Related Commands

Table 76 provides references to related commands.

<i>Table 76. Commands Related to LOCK ADMIN</i>		
Command	Description	Reference
UNLOCK ADMIN	Enables a locked administrator to access ADSM.	Page 435
QUERY ADMIN	Displays information about one or more ADSM administrators.	Page 256

LOCK NODE (Lock Out a Client Node)

Use the LOCK NODE command to prevent a client node from accessing the server.

By using the LOCK NODE command, you can restrict a client node from accessing the server. The client node is locked out until a system administrator uses the UNLOCK NODE command to reestablish access for the client node.

A locked client cannot perform backup and archive operations, regardless of whether the operations have been scheduled to occur.

Privilege Class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the client node belongs.

Syntax

```
▶▶ LOCK Node nodename ▶▶
```

Parameters

nodename

Specifies the name of the client node to be locked out. This parameter is required.

Examples

Task

Lock the client node TERRY.

Command: LOCK Node terry

Output:

```
ANR2064I Node TERRY locked.
```

Related Commands

Table 77 provides references to related commands.

Table 77. Commands Related to LOCK NODE

Command	Description	Reference
QUERY NODE	Displays partial or complete information about one or more clients.	Page 320
UNLOCK NODE	Enables a locked user in a specific policy domain to access the server.	Page 436

MACRO (Invoke a Macro)

Use the MACRO command to invoke a file that contains one or more ADSM administrative commands to be performed.

You cannot run macros from the server console. You can include MACRO commands within a macro to invoke other macros up to ten levels deep. A macro invoked from the batch or interactive mode of the administrative client is called a *high-level* macro. Any macros invoked from within the high-level macro are called *nested* macros.

Writing commands in a macro can be especially helpful when you want to issue commands that contain several parameters. You might also want to write macros to process related commands in a specific order. After you create a macro, you can update the information it contains and use it again, or you can copy the macro, make changes to the copy, then run the copy.

To document its purpose and function, add comments to your macro file by entering a slash and an asterisk (/), which identifies the beginning of the comment. Then enter the comment and an asterisk and a slash (*) to indicate the end of the comment. You can enter a comment on a line by itself, or enter it on a line that contains a command or part of a command. You cannot nest comments or span comments on more than one line. For more information, see "Using Macros to Issue Commands" on page 17.

If you have a series of commands that process successfully via the command-line, but are unsuccessful when issued within a macro, there are probably dependencies between commands. It is possible that a command issued within a macro cannot be processed successfully until a previous command, that is issued within the same macro, is committed. Either of the following actions allow successful processing of these commands within a macro:

- Insert a COMMIT command before the command dependent on a previous command. For example, if COMMAND C is dependent upon COMMAND B, you would insert a COMMIT command before COMMAND C. An example of this macro is:


```
COMMAND A  
COMMAND B  
COMMIT  
COMMAND C
```

- Start the administrative client session using the ITEMCOMMIT option. This causes each command within a macro to be committed before the next command is processed.

Privilege Class

Any administrator can issue this command.

Syntax



Parameters

macroname

Specifies the name of the macro. This parameter is required.

substitutionvalue

Specifies the value for a substitution variable in a macro. This parameter is optional. You can enter values for substitution variables to run a macro with substitution variables. When you use substitution variables, you can reuse a macro whenever you need to perform the same task for different objects or with different parameter values. Each substitution value must be delimited with a space. If you want to specify a value that contains blanks, you must enclose the value in quotation marks.

Examples

Task 1

Use the macro file REGNG to register and grant authority to a new administrator.

Macro Contents:

```
/* Register and grant authority to new administrator */
REGister Admin jones passwd -
CONtactinfo="x1235"
GRant AUTHority jones
CLasses=Policy
```

Command:MACRO regng

Output:

```
ANS5101I Server command: 'REGister Admin jones passwd
contact=x1235'
ANR2068I Administrator JONES registered.

ANS5101I Server command: 'GRant AUTHority jones
CLasses=policy'
ANR2077I Unrestricted policy privilege granted to
administrator JONES.
```

Task 2

Use the macro file AUTHRG, containing substitution variables, to register and grant authority to a new administrator.

When you run the macro, you then enter the values you want to pass to the server to process the command.

Macro Contents:

```
/* Registers and authorizes a new administrator */
REGister Admin %1 %2 - /* Enter userid and password */
CONTact=%3 /* Enter contact info (in quotes if nec.) */
GRant AUTHority %1 - /* Server uses variable already */
- /* defined by you */
CLasses=%4 /* Enter the privilege class */
```

Command:MACRO authrg jones passwd x1235 Policy

Output:

```
ANS5101I Server command: 'REGISTER ADMIN jones passwd
CONTACTinfo=x1235'
ANR2068I Administrator JONES registered.

ANS5101I Server command: 'GRANT AUTHority jones
CLasses=Policy'
ANR2077I Unrestricted policy privilege granted to
administrator JONES.
```

Related Commands

Table 78 provides references to related commands.

<i>Table 78. Commands Related to MACRO</i>		
Command	Description	Reference
COMMIT	Makes changes to the database permanent.	Page 68
ROLLBACK	Discards any uncommitted changes to the database since the last COMMIT was executed.	Page 405

MOVE DATA (Move Files on a Storage Pool Volume)

Use the MOVE DATA command to move data from either a primary or copy storage pool volume to other available storage pool volumes.

Restrictions:

1. You cannot move data from a primary storage pool volume to a copy storage pool.
2. You cannot move data on a copy storage pool volume to a different storage pool.

For additional information about storage pools, refer to *ADSM Administrator's Guide*.

Multiple sequential access volume mounts are required to move files that span volumes.

When moving files off of a random access volume, any cached copies of files on the volume (assuming caching is enabled) are erased. The MOVE DATA command is the only way to manually remove cached random access volume copies. (ADSM automatically removes cached files when the space they occupy is needed for primary files.)

If you are attempting to move files to volumes within the same storage pool, sufficient space must be available on the volumes or the operation fails.

You can use this command to move files from an offsite volume in a copy storage pool. Because the offsite volume cannot be mounted, the server obtains the active files on the offsite volume from either a primary storage pool or another copy storage pool. These files are then written to the destination volumes in the original copy storage pool.


This command can generate a background process that can be canceled with the CANCEL PROCESS command. If a MOVE DATA background process is canceled, some files may have already moved prior to the cancellation. To display information on background processes, use the QUERY PROCESS command.

Do not use the MOVE DATA command while a restore process (RESTORE STGPOOL or RESTORE VOLUME) is running. The MOVE DATA command could cause the restore to be incomplete.

Privilege Class

To issue this command, you must have system privilege, unrestricted storage privilege, or restricted storage privilege for the storage pool to which the volume belongs and also for the new storage pool, if one is specified.

Syntax



MOVE DATA

Parameters

volname

Specifies the name of the storage pool volume where the file to be moved currently resides. This parameter is required.

STGpool=*stgpoolname*

Specifies the primary storage pool name to which you want to move files from a primary storage pool volume. This parameter is optional and does not apply to copy storage pool volumes. If you do not specify a value for this parameter, files are moved to other volumes within the same storage pool.

Examples

Task 1

Move files from storage pool volume 201, assigned to storage pool POOL2, to any available volumes assigned to POOL2.

Command: MOVE Data 201

Output:

```
ANR2232W This command will move all of the files stored on
volume 201 to other volumes within the same storage pool;
the files will be inaccessible to users until the operation
completes.
Do you wish to proceed? (Y/N)
```

Type Y and press the ENTER key. A message similar to the following is displayed:

```
ANR1180I Move Data process started for volume 201
(process id 2).
ANS5104I Process number 2 started.
```

Task 2

Move files from storage pool volume HARVEY.ADSM.STG1, assigned to storage pool POOL2, to any available volumes assigned to the storage pool named BACKUPPOOL.

Command: MOVE Data harvey.adsm.stg1 STGpool=backuppool

Output:

```
ANR2232W This command will move all of the data stored on
volume HARVEY.ADSM.STG1 to other volumes in storage pool
BACKUPPOOL; the data will be inaccessible to users until the
operation completes.
Do you wish to proceed? (Y/N)
```

Type Y and press the ENTER key. A message similar to the following is displayed:

```

ANR1180I Move Data process started for volume
HARVEY.ADSM.STG1 (process id 2).
ANS5104I Process number 2 started.

```

Related Commands

Table 79 provides references to related commands.

Table 79. Commands Related to MOVE DATA

Command	Description	Reference
CANCEL PROCESS	Cancels a background server process.	Page 64
DEFINE VOLUME	Assigns a volume to be used for storage within a specified storage pool.	Page 148
DELETE VOLUME	Deletes a volume from a storage pool.	Page 185
QUERY PROCESS	Displays information about one or all background server processes.	Page 332

QUERY ACTLOG (Search Activity Log for Messages)

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.

Note: You cannot schedule the QUERY ACTLOG command using the DEFINE SCHEDULE command (central scheduling).

The activity log contains all messages that are sent to the server console under normal operation. These messages include session start and termination, background process messages, sequential access media mount information, and inventory expiration. Responses to commands entered at the server console are not recorded in the activity log unless the command affects or starts a background process or client session. In addition, all error messages are retained in the activity log.

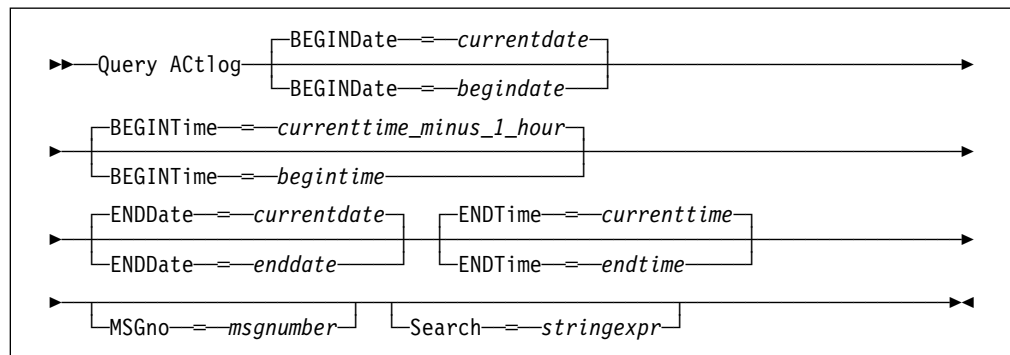
To minimize the processing time when querying the activity log, you should:

- Minimize the length of the time range for querying messages
- Restrict the message number or text string to those for which information is required

Privilege Class

Any administrator can issue this command.

Syntax



Parameters

BEGINdate=date

Specifies the beginning date of the range for the messages to be displayed, in MM/DD/YYYY format. This parameter is optional. All messages that occurred after this date are displayed. The default value is the current date. Possible values are:

MM/DD/YYYY

Specifies that you want to display information beginning with records created on a specific date, using MM/DD/YYYY format (for example, 01/31/1995).

TODAY

Specifies that you want to display information beginning with records created on the current date.

TODAY–numdays or –numdays

Specifies that you want to display information beginning with records created on the current date minus a specific number of days. The maximum value for the number of days is 9999.

For example, to display information beginning with records created a week ago, you can specify TODAY–7 or simply –7.

BEGINTime=currenttime_minus_1_hour

Specifies the beginning time of the range for messages to be displayed, in HH:MM format. This parameter is optional. All messages that occurred after this time are displayed. If you do not specify a value for this parameter, all messages that have occurred in the last hour are displayed.

ENDDate=date

Specifies the ending date of the range for messages to be displayed, in MM/DD/YYYY format. All messages that have occurred before this date are displayed. This parameter is optional. If you do not specify a value for this parameter, the current date is used. All the dates and times used are server date and time. Possible values are:

MM/DD/YYYY

Specifies that you want to display information ending with records created on a specific date using MM/DD/YYYY format (for example, 02/28/95).

TODAY

Specifies that you want to display information ending with records created on the current date.

TODAY–numdays or –numdays

Specifies that you want to display information ending with records created on or before the current date minus a specific number of days. The maximum value for the number of days is 9999.

For example, to display all volume history information created up to yesterday, you can specify TODAY–1 or simply –1.

ENDTime=currenttime

Specifies the ending time of the range for messages to be displayed, in HH:MM format. This parameter is optional. All messages that have occurred before this time are displayed. If you do not specify a value for this parameter, all messages are displayed up to the time that you issued this command.

MSGno=msgnumber

Specifies an integer that defines the number of the message to be displayed from the activity log. This is the numeric part of the message without the ANR prefix. This parameter is optional. If you do not specify a value for this parameter, all activity log messages are displayed.

Search=stringexpr

Specifies a string expression that you want to search for in the activity log. This parameter is optional. The string expression must be enclosed in quotation marks if it contains blanks or commas. You can use a pattern-matching

expression to specify this string. All matching strings are displayed. If you do not specify a value for this parameter, all strings are displayed.

Examples

Task 1

Search the activity log for all messages appearing on April 27, 1994 (04/27/94), between 14:57 and 15:39.

Command: Query ACTlog BEGINDate=04/27/1994 BEGINTime=14:57 ENDTIME=15:39

Output:

Date/Time	Message
04/27/1994 14:57:30	ANR0402I Session 17 started for administrator TOMC (0S/2).
04/27/1994 15:19:43	ANR0811I Inventory client file expiration started.
04/27/1994 15:19:43	ANR0812I Inventory client file expiration complete: 0 files deleted.
04/27/1994 15:31:30	ANR0482W Session 17 terminated - idle more than 15 minute(s).
04/27/1994 15:38:35	ANR0402I Session 18 started for administrator TOMC (0S/2).

where:

Date/Time

Specifies the date and time that the message was generated.

Message

Specifies the message.

Task 2

Search the activity log for any message that contains the string "delete." The output includes only messages produced during the past hour.

Command: Query ACTlog Search=delete

Output:

Date/Time	Message
04/27/1994 15:19:43	ANR0812I Inventory client file expiration complete: 0 files deleted.

where:

Date/Time

Specifies the date and time that the message was generated.

Message

Specifies the message.

Task 3

Display messages from the activity log that contain the text string "export." Search the activity log for messages that occurred on September 8, 1994 (09/08/1994) after 10:00.

Command:Query Actlog BEGINDate=09/08/1994 BEGINTime=10: Search=export*

Output:

Date/Time	Message
09-08-1994 10:00:27	ANR1360I Export volume EXP001 opened (sequence number 1).
09-08-1994 10:00:27	ANR0610I EXPORT SERVER started by SERVER_CONSOLE as process 4.
09-08-1994 10:00:27	ANR0639I EXPORT SERVER: Processing domain AIXDOMAIN.
09-08-1994 10:00:27	ANR0639I EXPORT SERVER: Processing domain STANDARD.
09-08-1994 10:00:27	ANR0640I EXPORT SERVER: Processing policy set ACTIVE in policy domain AIXDOMAIN.
09-08-1994 10:00:27	ANR0640I EXPORT SERVER: Processing policy set STANDARD in policy domain AIXDOMAIN.
09-08-1994 10:00:27	ANR0640I EXPORT SERVER: Processing policy set ACTIVE in policy domain STANDARD.
09-08-1994 10:00:27	ANR0640I EXPORT SERVER: Processing policy set STANDARD in policy domain STANDARD.
09-08-1994 10:00:27	ANR0641I EXPORT SERVER: Processing management class STANDARD in domain AIXDOMAIN, set ACTIVE.
09-08-1994 10:00:27	ANR0641I EXPORT SERVER: Processing management class STANDARD in domain AIXDOMAIN, set STANDARD.
09-08-1994 10:00:27	ANR0641I EXPORT SERVER: Processing management class STANDARD in domain STANDARD, set ACTIVE.
09-08-1994 10:00:27	ANR0641I EXPORT SERVER: Processing management class MC in domain STANDARD, set STANDARD.
09-08-1994 10:00:27	ANR0641I EXPORT SERVER: Processing management class STANDARD in domain STANDARD, set STANDARD.
09-08-1994 10:00:28	ANR0643I EXPORT SERVER: Processing archive copy group in domain AIXDOMAIN, set ACTIVE, management class STANDARD.

where:

Date/Time

Specifies the date and time that the message was generated.

Message

Specifies the message.

Related Commands

Table 80 provides a reference to a related command.

Table 80. Commands Related to QUERY ACTLOG

Command	Description	Reference
SET ACTLOGRETENTION	Specifies the number of days to retain log records in the activity log.	Page 409

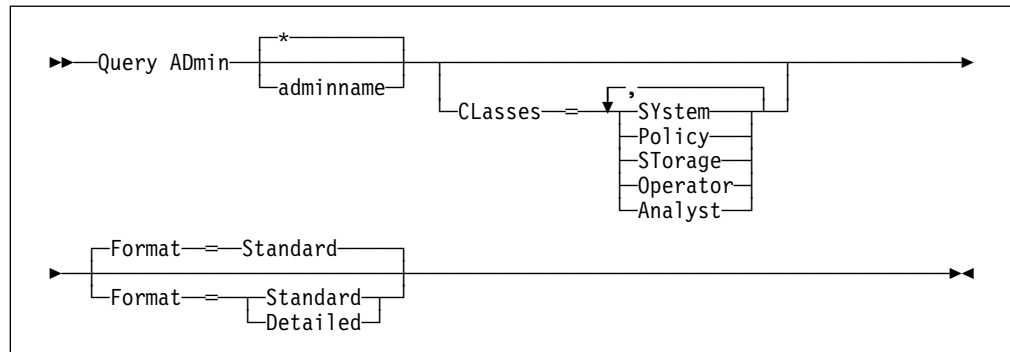
QUERY ADMIN (Display Information on One or More Administrators)

Use the QUERY ADMIN command to display information about one or more administrators.

Privilege Class

Any administrator can issue this command.

Syntax



Parameters

adminname

Specifies the name of the administrator for which you want to display information. This parameter is optional. You can use a pattern-matching expression to specify this name. All matching administrator names are included. If you do not specify a value for this parameter, all administrators are displayed.

Classes=classes

Specifies that you want to restrict output to those administrators which have been granted at least one of the specific privilege classes. This parameter is optional. You can specify more than one privilege class in a list, with each name delimited with a comma and no intervening spaces. If you do not specify a value for this parameter, information on all specified administrators is displayed, regardless of privilege class. Possible values are:

SSystem

Specifies that you want to display information on system administrators. A system administrator has the highest level of authority in ADSM. A system administrator can issue any administrative command and has authority to manage all policy domains and all storage pools. Only a system administrator can grant authority to other administrators.

Policy

Specifies that you want to display information on policy administrators (unrestricted and restricted). A policy administrator can issue commands that relate to policy management objects such as policy domains, policy sets, management classes, copy groups, and schedules.

STorage

Specifies that you want to display information on storage administrators (unrestricted and restricted). A storage administrator can issue commands that allocate and control storage resources for the server.

Operator

Specifies that you want to display information on operator administrators. An administrator with operator privilege can issue commands that affect the immediate operation of the server, including server availability and physical media.

Analyst

Specifies that you want to display information on analyst administrators. An administrator with analyst privilege can issue commands that reset the counters which track server performance.

Format=*formatvalue*

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed for the specified administrators.

Detailed

Specifies that complete information is displayed for the specified administrators.

Examples

Task 1

Display partial information on all administrators.

Command: Query Admin

Output:

Administrator Name	Days Since Last Access	Days Since Password Set	Locked?	Privilege Classes
SHELLBEE	3	21	No	Policy Operator
MIKE	11	11	No	Policy
SERVER_CONSOLE			No	System

where:

Administrator Name

Specifies the name of the administrator.

Days Since Last Access

Specifies the number of days since the administrator last accessed the server.

Days Since Password Set

Specifies the number of days since the administrator's password was defined or most recently updated.

Locked?

Specifies whether the administrator is locked out of the system.

Privilege Classes

Specifies the privilege classes for which the administrator has authority.

Task 2

Display complete information for the administrator named MIKE.

Command: Query Admin Format=Detailed

Output:

```
Administrator Name: MIKE
Last Access Date/Time: 03/31/1994 17:59:07
Days Since Last Access: 1
Password Set Date/Time: 03/31/1994 17:59:07
Days Since Password Set: 1
Locked?: No
Contact:
System Privilege:
Policy Privilege: EMPLOYEE_RECORDS PROG1 STANDARD
Storage Privilege:
Analyst Privilege:
Operator Privilege: Yes
Registration Date/Time: 03/31/1994 17:59:07
Registering Administrator: RAY
```

where:

Administrator Name

Specifies the name of the administrator.

Last Access Date/Time

Specifies the date and time that the administrator last accessed the server.

Days Since Last Access

Specifies the number of days since the administrator last accessed the server.

Password Set Date/Time

Specifies the date and time that the administrator's password was defined or most recently updated.

Days Since Password Set

Specifies the number of days since the administrator's password was defined or most recently updated.

Locked?

Specifies whether the administrator is locked out of the system.

Contact

Specifies any contact information for the administrator.

System Privilege

Specifies whether the administrator has been granted system privilege.

Policy Privilege

Specifies whether the administrator has been granted unrestricted policy privilege or the names of any policy domains that the restricted policy administrator can manage.

Storage Privilege

Specifies whether the administrator has been granted unrestricted storage privilege or the names of any storage pools that the restricted storage administrator can manage.

Analyst Privilege

Specifies whether the administrator has been granted analyst privilege.

Operator Privilege

Specifies whether the administrator has been granted operator privilege.

Registration Date/Time

Specifies the date and time that the administrator was registered.

Registering Administrator

Specifies the name of the administrator who registered the administrator.

Related Commands

Table 81 provides references to related commands.

Table 81. Commands Related to QUERY ADMIN

Command	Description	Reference
GRANT AUTHORITY	Assigns one or more privilege classes to an administrator.	Page 217
REGISTER ADMIN	Defines a new administrator without granting administrative authority.	Page 370
REMOVE ADMIN	Removes an administrator from the list of registered administrators.	Page 377
RENAME ADMIN	Changes an ADSM administrator's name.	Page 380
REVOKE AUTHORITY	Revokes one or more privilege classes or restricts access to policy domains and storage pools.	Page 401


QUERY ASSOCIATION (Query Client Node Associations with a Schedule)

Use the QUERY ASSOCIATION command to display the client nodes associated with one or more schedules. Client nodes that are associated with a schedule perform backup or archive operations according to that schedule.

Privilege Class

Any administrator can issue this command.

Syntax



Parameters

domainname

Specifies the name of the policy domain that you want to display. You can use a pattern-matching expression to specify this name. All matching policy domain names are displayed. If you do not specify a value for this parameter, all existing policy domains are queried. If you specify a domain name, you do not have to specify a schedule name.

schedulingname

Specifies the name of the schedule that you want to display. You can use a pattern-matching expression to specify this name. All matching schedule names are displayed. If you do not specify a value for this parameter, all existing schedules are queried. If you specify a schedule name, you must also specify a policy domain name.

Examples

Task 1

Display all the clients that are associated with each schedule that belongs to the OS2_DOMAIN policy domain.

Command: Query ASSOCIation os2_domain *

Output:

```
Policy Domain Name: OS2_DOMAIN
Schedule Name: WINTER
Associated Nodes: TMCQ DAVE MARIE CLAUDIA DAVID

Policy Domain Name: OS2_DOMAIN
Schedule Name: WEEKLY_BACKUP
Associated Nodes: BOB FR̄ED GEORGE CANNON BONZO KACZ
DELANEY SUSAN MARY FRANK JONES ANDERSON
NANCY NODE5 JIM KAREN NODE22
```

where:

Policy Domain Name

Specifies the name of the policy domain to which the schedule belongs.

Schedule Name

Specifies the name of the schedule that belongs to the specified policy domain.

Associated Nodes

Specifies the names of the clients that are currently associated with the specified schedule.

Task 2

Display all the clients that are assigned to the DOMAIN1 policy domain and associated with the WINTER schedule.

Command: Query ASSOCIATION domain1 winter

Output:

```
Policy Domain Name: DOMAIN1
Schedule Name: WINTER
Associated Nodes: DAVE DAVID LARRY MARIE
```

where:

Policy Domain Name

Specifies the name of the policy domain to which the schedule belongs.

Schedule Name

Specifies the name of the schedule that belongs to the specified policy domain.

Associated Nodes

Specifies the names of the clients that are currently associated with the specified schedule.

Related Commands

Table 82 provides references to related commands.

<i>Table 82. Commands Related to QUERY ASSOCIATION</i>		
Command	Description	Reference
DEFINE ASSOCIATION	Associates one or more clients with a schedule.	Page 81
DELETE ASSOCIATION	Deletes the association between one or more clients and a schedule.	Page 151

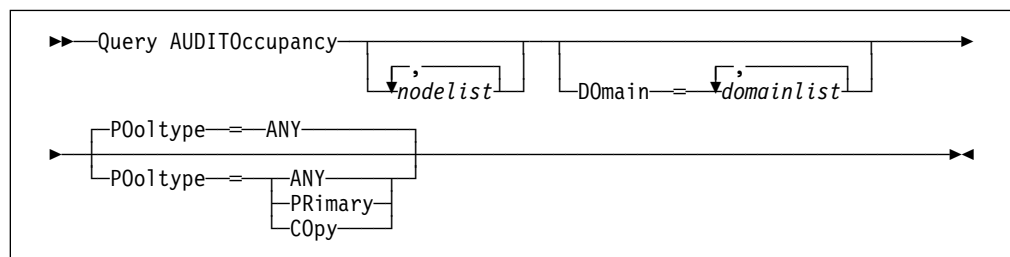
QUERY AUDITOCUPANCY (Query Client Node Storage Utilization)

Use the QUERY AUDITOCUPANCY command to display information about the client node data 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 storage usage in primary and copy storage pools, by client node. In addition, you can use this information to determine if and where client node storage usage needs to be balanced. This information can also assist with billing clients for storage usage.

Privilege Class

Any administrator can issue this command.

Syntax



Parameters

nodelist

Specifies a list of nodes for which data storage utilization information is to be displayed. You can specify more than one node by separating the node names using commas, with no intervening spaces. This parameter is optional. You can use pattern-matching expressions to specify names. The default (*) is to query all client nodes. You can also limit this list by policy domain via the optional DOMAIN parameter.

Domain=*domainlist*

Specifies a list of policy domains to restrict which nodes are displayed. This parameter is optional. Only nodes belonging to the specified policy domains are displayed. You can specify more than one policy domain by separating the policy domain names with commas and no intervening spaces. You can use pattern-matching expressions to specify names.

Pooltype=*value*

Specifies whether to display storage usage in primary storage pools only, copy storage pools only, or both primary and copy storage pools. This parameter is optional. The default is ANY. Possible values are:

ANY

Specifies that the storage usage displayed includes files stored in both primary and copy storage pools. The value presented is the total for the two pools.

PRimary

Specifies that storage usage in primary storage pools is displayed.

COPY

Specifies that storage usage in copy storage pools is displayed.

Examples

Task

Display combined storage utilization in primary and copy storage pools.

Command: Query AUDITOccupancy

Output:

```

License information as of last audit on 03/04/1995 at 08:43:57.

Client Node Name      Backup      Archive      Space-Managed      Total
                       Storage      Storage      Storage Used        Storage
                       Used (MB)    Used (MB)    (MB) Used (MB)
-----
POLICY                245         20           0                   265
FRED                  300         15           5                   320
    
```

where:

Client Node Name

Specifies the name of the client node.

Backup Storage Used (MB)

Specifies the total backup storage utilization for the node.

Archive Storage Used (MB)

Specifies the total archive storage utilization for the node.

Space-Managed Storage Used (MB)

Specifies the amount of server storage used to store files migrated from the client node.

Total Storage Used (MB)

Specifies the total client storage utilization (for backup, archive, and space management) for the node.

Related Commands

Table 83 provides references to related commands.

Table 83 (Page 1 of 2). Commands Related to QUERY AUDITOCUPANCY

Command	Description	Reference
AUDIT LICENSES	Checks for compliance with defined licenses.	Page 43
QUERY LICENSE	Displays information about licenses and audits.	Page 304

QUERY AUDITOCUPANCY

Table 83 (Page 2 of 2). Commands Related to QUERY AUDITOCUPANCY

Command	Description	Reference
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.	Page 342
REGISTER LICENSE	Registers a new license with the ADSM server.	Page 372
SET LICENSEAUDITPERIOD	Specifies the number of days between automatic license audits.	Page 414

QUERY CONTENT (Query the Contents of a Storage Pool Volume)

Use the QUERY CONTENT command to display information about one or more files currently residing in a storage pool volume. This command can also be used to identify files that have been found to be damaged and files that have been backed up to a copy storage pool.

This command is useful when a volume is damaged or before you:

- Request the server to fix any inconsistencies between a volume and the database
- Move files from one volume to another volume
- Delete a volume from a storage pool

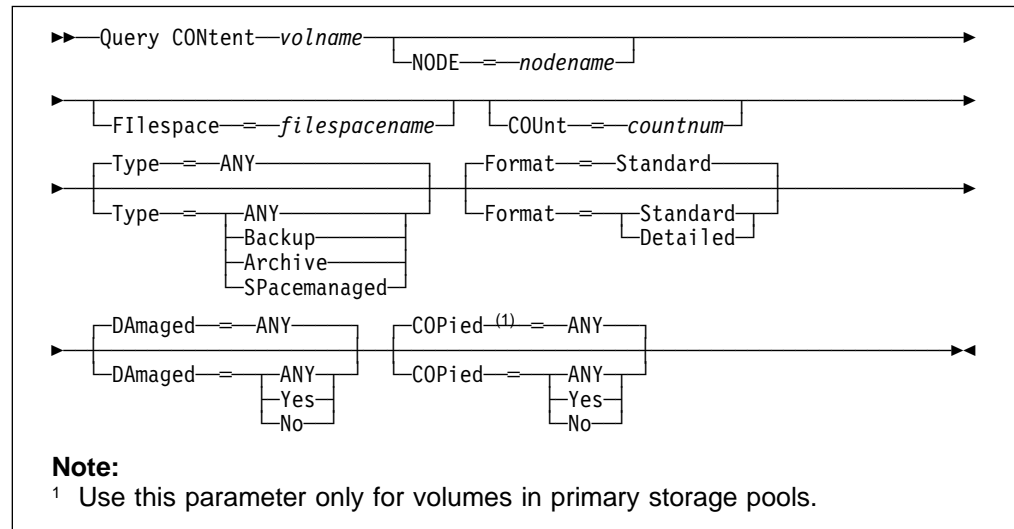
Because the results of this command can take a long time to execute and can be extremely large, it is recommended that you use the COUNT parameter to limit the number of files displayed.

Note: Files in cache that are marked as damaged are never displayed.

Privilege Class

Any administrator can issue this command.

Syntax



Parameters

volname

Specifies the name of the volume to be queried. This parameter is required.

NODE=*nodename*

Specifies the name of the backup-archive client or space management client associated with the file space to be queried. This parameter is optional. You can use a pattern-matching expression to specify this name. If you do not specify a name, all backup-archive and space management clients are included.

Filespace=*filespace**name*

Specifies the name of the file space to queried. This parameter is optional. You can use a pattern-matching expression to specify this name. File space names are case-sensitive. If you do not specify a file space name, all file spaces are included.

COUnt=*count**num*

Specifies the number of files to be displayed. This parameter is optional. You can specify either a negative integer or a positive integer. If you specify a positive integer, *n*, the first *n* files are displayed. If you specify a negative integer, *-n*, the last *n* files are displayed in *reverse* order. You cannot specify COUNT=0. If you do not specify a value for this parameter, all files are displayed.

Type=*file**type*

Specifies the type of file space to be queried. This parameter is optional. The default value is ANY. Possible values are:

ANY

Specifies that all types of files in the specified storage pool volume are queried: backup versions of files, archived copies of files, and files migrated from one or more client nodes.

Backup

Specifies that only backup file spaces are queried.

Archive

Specifies that only archive file spaces are queried.

SPacemanaged

Specifies that only space management file spaces (files migrated from one or more client nodes) are queried.

Format=*format**value*

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed for the specified storage pool volume.

Detailed

Specifies that complete information is displayed for the specified storage pool volume.

DAmaged=*value*

Specifies whether to restrict the query output either to files that are known to be damaged, or to files that are not known to be damaged. This parameter is optional. The default value is ANY. Possible values are:

ANY

Specifies that files are displayed regardless of whether the files have been found to be damaged.

Yes

Specifies that only files that have been damaged are displayed. These are files that were found to have data integrity errors when a user

attempted to restore, retrieve, or recall the file, or when an AUDIT VOLUME operation was run.

No

Specifies that only files *not* known to have data integrity errors are displayed.

COPIED=*value*

Specifies whether to restrict the query output either to files that are backed up to a copy storage pool, or to files that are not backed up to a copy storage pool. This parameter is optional. The default value is ANY. Possible values are:

ANY

Specifies that files are displayed regardless of whether the files are backed up to a copy storage pool. Primary and cached file copies are displayed.

Yes

Specifies that the only files displayed are those for which at least one, usable backup copy exists in a copy storage pool. A file is not displayed if its copy in the copy storage pool is known to have data integrity errors. Cached file copies are not displayed, because these files are never restored.

Use COPIED=YES to identify primary files that can be restored using the RESTORE VOLUME or RESTORE STGPOOL command.

No

Specifies that the only files displayed are those for which no usable backup copies exist in a copy storage pool. Cached file copies are not displayed, because these files are never restored.

Use COPIED=NO to identify primary files that cannot be restored using the RESTORE VOLUME or RESTORE STGPOOL command.

Examples

Task 1

Query the contents of the volume named JENVO5. Only display information about files migrated from client node PEGASUS.

Command: Query CONTENT jenvo5 Type=SPacemanaged

Output:

Node Name	Type	Filespace Name	Client's Name for File
PEGASUS	SpMg	/migtest	/jen

where:

Node Name

Specifies the name of the node to which the specified file belongs.

QUERY CONTENT

Type

Specifies the type of file: backup, archive, or space managed.

Filespace Name

Specifies the name of the file space to which the specified file belongs.

Client's Name for File

Specifies the name of the file.

Task 2

Query the contents of the volume named 200. Only display the last 3 files of the output.

Command: Query CONtent 200 Count=-3 Format=Detailed

Output:

```
Node Name: CASE6K
Type: Bkup
Filespace Name : /home
Client's Name for File: /case/ d.d
Stored Size: 84
Segment Number: 1/1
Cached Copy?: No

Node Name: CASE6K
Type: Bkup
Filespace Name : /home
Client's Name for File: /case/ c.c
Stored Size: 93
Segment Number: 1/1
Cached Copy?: No

Node Name: CASE6K
Type: Bkup
Filespace Name : /home
Client's Name for File: /case/ a.a
Stored Size: 93
Segment Number: 1/1
Cached Copy?: No
```

where:

Node Name

Specifies the name of the node to which the specified file belongs.

Type

Specifies the type of file, either BACKUP or ARCHIVE.

Filespace Name

Specifies the name of the file space to which the specified file belongs.

Client's Name for File

Specifies the name of the file.

Stored Size

Specifies the size of the file, in bytes.

Segment Number

Specifies whether the file is stored across multiple volumes. For example, if the specified file is stored across two volumes, the segment number indicates 1/2 (the first part of the file is stored on the specified volume) or 2/2 (the second part of the file is stored on the specified volume). If the segment number is 1/1, the file is completely stored on the specified volume.

Cached Copy?

Specifies whether the specified file is a cached copy of a migrated file.

Related Commands

Table 84 provides references to related commands.

Table 84. Commands Related to QUERY CONTENT

Command	Description	Reference
BACKUP STGPOOL	Backs up a primary storage pool to a copy storage pool.	Page 56
DEFINE VOLUME	Assigns a volume to be used for storage within a specified storage pool.	Page 148
DELETE VOLUME	Deletes a volume from a storage pool.	Page 185
RESTORE STGPOOL	Restores files to a primary storage pool from copy storage pools.	Page 391
RESTORE VOLUME	Restores files stored on specified volumes in a primary storage pool from copy storage pools.	Page 396
UPDATE VOLUME	Updates the attributes of one or more storage pool volumes.	Page 495

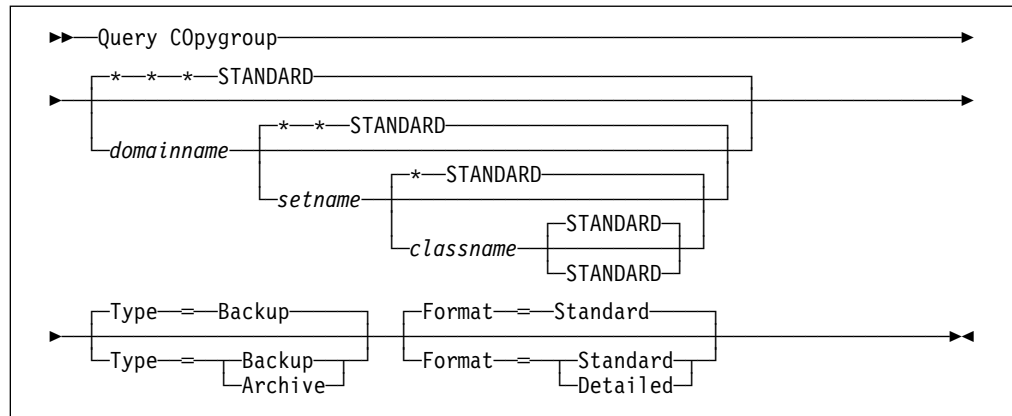
QUERY COPYGROUP (Query One or More Copy Groups)

Use the QUERY COPYGROUP command to display information about one or more copy groups.

Privilege Class

Any administrator can issue this command.

Syntax



Parameters

domainname

Specifies the name of the policy domain associated with the copy group to be queried. This parameter is optional. You can use a pattern-matching expression to specify this name. All matching policy domains are queried. If you do not specify a value for this parameter, all policy domains are queried. You must specify this parameter when querying an explicitly named copy group.

setname

Specifies the name of the policy set associated with the copy group to be queried. This parameter is optional. You can use a pattern-matching expression to specify this name. All matching policy sets are queried. If you do not specify a value for this parameter, all policy sets are queried. You must specify this parameter when querying an explicitly named copy group.

classname

Specifies the name of the management class associated with the copy group to be queried. This parameter is optional. You can use a pattern-matching expression to specify this name. All matching management classes are queried. If you do not specify a value for this parameter, all management classes are queried. You must specify this parameter when querying an explicitly named copy group.

STANDARD

Specifies the name of the copy group. This parameter is optional. The name of the copy group must be STANDARD. If you do not specify a value for this parameter, the system defaults to STANDARD.

Format=*formatvalue*

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed for the specified copy group.

Detailed

Specifies that complete information is displayed for the specified copy group.

Type=*type*

Specifies the type of copy group to be queried. This parameter is optional. The default value is BACKUP. Possible values are:

Backup

Specifies that you want to query backup copy groups.

Archive

Specifies that you want to query archive copy groups.

Examples

Task 1

Display complete information on the backup copy group assigned to the MCLASS1 management class in the ACTIVE policy set of the DOMAIN3 policy domain.

Command: Query C0pygroup domain3 active mclass1 Format=Detailed

Output:

```

Policy Domain Name: DOMAIN3
Policy Set Name: ACTIVE
Mgmt Class Name: MCLASS1
Copy Group Name: STANDARD
Copy Group Type: Backup
Versions Data Exists: 1
Versions Data Deleted: 1
Retain Extra Versions: 30
Retain Only Version: 60
Copy Mode: Modified
Copy Serialization: 1
Copy Frequency: 1
Copy Destination: 30
Last Update by (administrator): TOMC
Last Update Date/Time: 04/27/1994 10:30:07
    
```

See Task 2 for parameter definitions.

Task 2

Display complete information on the archive copy group STANDARD that is assigned to the MCLASS1 management class in the SUMMER policy set of the PROG1 policy domain.

Command: Query C0pygroup prog1 summer mclass1 Type=Archive
Format=Detailed

QUERY COPYGROUP

Output:

```
Policy Domain Name: PROG1
Policy Set Name: SUMMER
Mgmt Class Name: MCLASS1
Copy Group Name: STANDARD
Copy Group Type: Archive
Retain Version: 730
Copy Serialization: Shared Static
Copy Frequency: CMD
Copy Mode: Absolute
Copy Destination: ARCHPOOL
Last Update by (administrator): SERVER_CONSOLE
Last Update Date/Time: 04/15/1994 09:26:20
```

where:

Policy Domain Name

Specifies the name of the policy domain.

Policy Set Name

Specifies the name of the policy set.

Mgmt Class Name

Specifies the name of the management class.

Copy Group Name

Specifies the name of the copy group. This name must be STANDARD.

Copy Group Type

Specifies the type of the copy group.

Versions Data Exists

Specifies the maximum number of backup versions that are retained by the server for files that are on the client's file system.

Versions Data Deleted

Specifies the maximum number of backup versions that are retained by the server for files that have been deleted from the client's file system.

Retain Version

Specifies the number of days to keep an archive copy of a file bound to the management class.

Retain Extra Versions

Specifies the amount of time to retain all backup versions, except for the most recent backup copy.

Retain Only Version

Specifies the retention time, in days, for the last backup version of a file that has been deleted from the client's file system.

Copy Serialization

Specifies whether a file can be in use during an archive operation.

Copy Frequency

Specifies the copy frequency of the copy group. For archive copy groups, this value must be CMD.

Copy Mode

Specifies that files in the copy group are archived regardless of whether they have been modified. For archive copy groups, this value must be ABSOLUTE.

Copy Destination

Specifies the name of the storage pool where the archive copy group's files are to reside.

Last Update by (administrator)

Specifies the name of the administrator that defined or most recently updated the copy group.

Last Update Date/Time

Specifies the date and time that the administrator defined or most recently updated the copy group.

Related Commands

Table 85 provides references to related commands.

Table 85. Commands Related to QUERY COPYGROUP

Command	Description	Reference
DEFINE COPYGROUP	Defines a copy group for backup or archive processing within a specified management class.	Page 83
DELETE COPYGROUP	Deletes a backup or archive copy group from a policy domain and policy set.	Page 153
UPDATE COPYGROUP	Changes one or more attributes of a copy group.	Page 439


QUERY DB (Display Information on the Database)

Use the QUERY DB command to display allocation information for the database, including information about utilization, expansion, and reduction abilities.

Privilege Class

Any administrator can issue this command.

Syntax



Parameters

Format=*formatvalue*

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed for the database.

Detailed

Specifies that complete information is displayed for the database.

Examples

Task 1

Display statistical information about the database.

Command: Query DB

Output:

Available Space (MB)	Assigned Capacity (MB)	Maximum Extension (MB)	Maximum Reduction (MB)	Page Size (bytes)	Total Pages	Used Pages	%Util	Max. %Util
8	8	0	0	4,096	2,048	1,377	67.2	67.2

See Task 2 for parameter definitions.

Task 2

Display statistical information about the database in complete detail.

Command: Query DB Format=Detailed

Output:

```

    Available Space (MB): 396
    Assigned Capacity (MB): 396
    Maximum Extension (MB): 0
    Maximum Reduction (MB): 376
        Page Size (bytes): 4,096
    Total Usable Pages: 101,376
        Used Pages: 4,151
            %Util: 4.1
        Max. %Util: 4.1
        Changed Pages: 1
        Physical Volumes: 1
        Buffer Pool Pages: 128
    Total Buffer Requests: 63,264
        Cache Hit Pct.: 96.25
        Cache Wait Pct.: 0.00
    Backup in Progress?: No
    Type of Backup in Progress:
    Incrementals Since Last Full: 12
    Changed Since Last Backup (MB): 0.25
        Percentage Changed: 1.54
    Last Complete Backup Date/Time: 03/01/1995 12:07:15

```

where:

Available Space (MB)

Specifies the amount of space available in the database, in megabytes.

Assigned Capacity (MB)

Specifies the assigned storage capacity for the database, in megabytes.

Maximum Extension (MB)

Specifies the maximum amount the database can be extended, in megabytes.

Maximum Reduction (MB)

Specifies the maximum amount the database can be reduced, in megabytes.

Page Size (bytes)

Specifies the page size, in bytes.

Total Usable Pages

Specifies the total number of pages available in the database. A page is a unit of allocation within database volumes. The total number of pages will always be a multiple of the Assigned Capacity value, by a factor of 256 (1 megabyte/256 pages).

Used Pages

Specifies the total number of pages currently being used by the database.

%Util

Specifies the current utilization as a percentage value of the assigned capacity.

Max. %Util

Specifies the maximum occupancy for the database, as a percentage of the available space. This value specifies the highest capacity occupied, or "high water mark" for database utilization. You can reset this value with the RESET DBMAXUTILIZATION command.

Changed Pages

Specifies the total number of pages changed since the last backup of the database.

Physical Volumes

Specifies the number of volumes defined to the database, including database copies.

Buffer Pool Pages

Specifies the number of pages in the database buffer pool. This value is determined by the value for the BUFFPOOLSIZE option. The default value for the BUFFPOOLSIZE option is 512 kilobytes (128 pages).

Total Buffer Requests

Specifies the cumulative number of requests for database pages, since the last RESET BUFPOOL command (or since the server was last started). You can reset this value with the RESET BUFPOOL command.

Cache Hit Pct.

Specifies, as a percentage of the total number of database requests, the number of requests for a database page that is already in the database buffer pool. Requesting pages that are already in the buffer pool increases database performance. You can reset this value with the RESET BUFPOOL command.

Cache Wait Pct.

Specifies, as a percentage, the number requests for a database buffer pool page that was unavailable (because all database buffer pool pages are occupied). If this number is greater than 0.0, you should increase the size of the buffer pool with the BUFFPOOLSIZE option. You can reset this value with the RESET BUFPOOL command.

Backup in Progress?

Specifies whether a backup of the database is in progress.

Type of Backup in Progress

Specifies whether the backup in progress is a full or incremental backup.

Incremental Since Last Full

Specifies the number of incremental backups that have been run of the database since the last full backup was done.

Changed Since Last Backup (MB)

Specifies the pages (in megabytes) that have changed since the last backup of the database.

Percentage Changed

Specifies, as a percentage, the number of pages that changed since the last backup of the database.

Last Complete Backup Date/Time

Specifies the data and time of the last complete backup of the database.

Related Commands

Table 86 provides references to related commands.

Table 86. Commands Related to QUERY DB

Command	Description	Reference
BACKUP DB	Backs up the ADSM database to sequential access volumes.	Page 50
DEFINE DBCOPY	Creates a volume copy of a database volume.	Page 95
DEFINE DBVOLUME	Assigns a volume to be used for database storage.	Page 97
EXTEND DB	Extends usable database storage within the volumes previously allocated for its use.	Page 213
EXTEND LOG	Extends usable log storage within the volumes previously allocated for its use.	Page 215
REDUCE DB	Attempts to free up a database volume by reducing usable storage within the volumes defined.	Page 366
REDUCE LOG	Attempts to free up a log volume by reducing usable storage within the volumes defined.	Page 368
RESET BUFPOOL	Resets the buffer pool statistics for the database.	Page 386
RESET DBMAXUTILIZATION	Resets the maximum utilization statistic for the database.	Page 388


QUERY DBBACKUPTRIGGER (Query the Database Backup Trigger)

Use the QUERY DBBACKUPTRIGGER command to display the current settings for the database backup trigger.

Privilege Class

Any administrator can issue this command.

Syntax



Parameters

Format=*formatvalue*

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed for the database backup trigger.

Detailed

Specifies that complete information is displayed for the database backup trigger.

Examples

Task

Display the current settings for the database backup trigger.

Command: Query DBBackuptrigger Format=Detailed

Output:

```

    Full Device Class: FILE
    Incremental Device Class: FILE
    Log Full Percentage: 50
    Incrementals Between Fulls: 6
    Last Update by (administrator): SERVER_CONSOLE
    Last Update Date/Time: 05/31/1995 11:15:35
  
```

where:

Full Device Class

Specifies the name of the device class used for full backups.

Incremental Device Class

Specifies the name of the device class used for incremental backups.

Log Full Percentage

Specifies the percentage of space used for the recovery log at which ADSM performs an automatic backup of the database.

Incrementals Between Fulls

Specifies the number of incremental backups ADSM runs between full backups of the database.

Last Update by (administrator)

Specifies the name of the administrator that most recently updated the database.

Last Update Date/Time

Specifies when the administrator most recently updated the database.

Related Commands

Table 87 provides references to related commands.

Table 87. Commands Related to QUERY DBBACKUPTRIGGER

Command	Description	Reference
DEFINE DBBACKUPTRIGGER	Defines a trigger for automatic ADSM database backups.	Page 91
DELETE DBBACKUPTRIGGER	Deletes the database backup trigger.	Page 156
UPDATE DBBACKUPTRIGGER	Changes one or more attributes of the database backup trigger.	Page 447


QUERY DBVOLUME (Display Information on One or More Database Volumes)

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.

Privilege Class

Any administrator can issue this command.

Syntax



Parameters

volname

Specifies the name of the database volume to be queried. This parameter is optional. You can use a pattern-matching expression to specify this name. All matching volume names are displayed. If you do not specify a volume name, information on all database volumes is displayed.

Format=*formatvalue*

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information for the specified database volumes is displayed.

Detailed

Specifies that complete information for the specified database volumes is displayed.

Examples

Task

Display information in complete detail on the database volume at address HARVEY.ADSM.PINV.

Command: Query DBVolume harvey.adsm.pinv Format=Detailed

Output:

```

Volume Name (Copy 1): HARVEY.ADSM.PINV
      Copy Status: Sync'd
Volume Name (Copy 2):
      Copy Status: Undefined
Volume Name (Copy 3):
      Copy Status: Undefined
Available Space (MB): 8
Allocated Space (MB): 8
      Free Space (MB): 0
    
```

where:

Volume Name (Copy *n*)

Specifies the name of the database volume. The output displays status of the three volume copies (Copy 1, Copy 2, and Copy 3).

Copy Status

Specifies the status of the database volume. Possible values are:

Offline

The volume is varied offline.

Stale

The volume is not synchronized with the other volumes in the database and cannot be used for storage.

Sync'd

The volume is synchronized with the other volumes in the database and can be used for storage.

Undefined

No volume has been defined.

Available Space (MB)

Specifies the total amount of space that is available, both allocated and free, for the database volume, in megabytes.

Allocated Space

Specifies the amount of space allocated to the database volume, in megabytes.

Free Space

Specifies the amount of unassigned space for this mirror set that can be allocated for use, in megabytes.

Related Commands

Table 88 provides references to related commands.

Table 88. Commands Related to QUERY DBVOLUME

Command	Description	Reference
DEFINE DBCOPY	Creates a volume copy of a database volume.	Page 95
DEFINE DBVOLUME	Assigns a volume to be used for database storage.	Page 97
DELETE DBVOLUME	Deletes a volume from the database.	Page 158
VARY	Specifies whether a disk volume is available to the server for use.	Page 504


QUERY DEVCLASS (Display Information on One or More Device Classes)

Use the QUERY DEVCLASS command to display information on one or more device classes.

Privilege Class

Any administrator can issue this command.

Syntax



Parameters

devclassname

Specifies the name of the device class to be queried. This parameter is optional. You can use wildcard characters to specify this name. All matching device classes are displayed. If you do not specify a value for this parameter, all device classes are displayed.

Format=*formatvalue*

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed for the specified device class.

Detailed

Specifies that complete information is displayed for the specified device class.

Usage Notes

- The mount wait value is now displayed for a device class on an MVS server.

Examples

Task 1

Display information on the device class name SQUARE1.

Command: Query DEVclass square1

Output:

Device Class Name	Device Access Strategy	Storage Device Pool Count	Device Type	Estimated Capacity (MB)	Mount Limit	Retention (min)	Unit Name	Comp
SQUARE1	Sequential	1	Cartridge	180.0	2	60	3480	Yes

QUERY DEVCLASS

where:

Device Class Name

Specifies the name of the device class.

Device Access Strategy

Specifies how data is written to the device class.

Storage Pool Count

Specifies the number of storage pools that are assigned to the device class.

Device Type

Specifies the device type of the device class.

Estimated Capacity (MB)

Specifies the estimated capacity of a volume associated with this device class.

Mount Limit

Specifies the maximum number of sequential access volumes that can be mounted concurrently.

Mount Retention (min)

Specifies the maximum amount of time, in minutes, to keep an *idle* sequential access volume mounted.

Unit Name

Specifies the esoteric unit name of the device for allocation. The default is 3480.

Comp

Specifies whether compression is enabled for the device.

Task 2

Display, in full detail, information on the device class named SQUARE2.

Command: Query DEVclass square2 Format=Detailed

Output:

```
Device Class Name: SQUARE2
Device Access Strategy: Sequential
Storage Pool Count: 0
Last Update by (administrator):
Last Update Date/Time: 05/15/1995 18:45:19
Device Type: Cartridge
Maximum Capacity (MB):
Estimated Capacity (MB): 180.0
Dataset Name Prefix: ADSM
Mount Limit: 2
Mount Retention (min): 60
Label Type: IBMSL
Expiration Date:
Mount Wait (min): 60
Unit Name: 3480
Compression: Yes
Protection: No
Retention:
```

where:

Device Class Name

Specifies the name of the device class.

Device Access Strategy

Specifies how data is written to the device class.

Storage Pool Count

Specifies the number of storage pools that are assigned to the device class.

Last Update by (administrator)

Specifies the name of the administrator that last modified the device class.

Last Update Date/Time

Specifies the time and date that the device class was last modified.

Device Type

Specifies the device type of the device class.

Maximum Capacity (MB)

Specifies the maximum capacity of a volume associated with this device class.

Estimated Capacity (MB)

Specifies the estimated capacity of a volume associated with this device class.

Dataset Name Prefix

Specifies the prefix for the device class.

Mount Limit

Specifies the maximum number of sequential access volumes that can be mounted.

Mount Retention (min)

Specifies the amount of time, in minutes, to keep an *idle* sequential access volume mounted.

Label Type

Specifies the label type as IBM standard labels.

Expiration Date

Specifies the expiration date placed on the sequential access label which is used for allocation.

Mount Wait (min)

Specifies the amount of time, in minutes, that the server waits for a sequential access volume mount request to be satisfied before canceling the request.

Unit Name

Specifies the esoteric unit name of the device class for allocation. The default is 3480.

Compression

Specifies whether compression is enabled for the device.

Protection

Specifies whether RACF protection is available for the device.

QUERY DEVCLASS

Retention

Specifies the number of days to retain data assigned to this device.

Related Commands

Table 89 provides references to related commands.

Table 89. Commands Related to QUERY DEVCLASS

Command	Description	Reference
DEFINE DEVCLASS	Defines a device class that can be used for storage pools, importing, or exporting.	Page 99
DELETE DEVCLASS	Deletes a device class name.	Page 160
UPDATE DEVCLASS	Changes the attributes of a device class.	Page 450

QUERY DOMAIN (Query One or More Policy Domains)


Use the QUERY DOMAIN command to display information on one or more policy domains. The QUERY DOMAIN command is useful in displaying the name of the policy set last activated and its default management class.

There is a difference between the “activated policy set,” as displayed in the QUERY DOMAIN command, and the “ACTIVE” policy set. When a policy set is activated, the contents of the policy set are copied to a policy set that has the reserved name ACTIVE. Once activated, there is no real relationship between the policy set that has been activated (copied to ACTIVE) and the contents of the ACTIVE policy set. The original policy set can still be modified, but the copied definitions in the ACTIVE policy set can only be modified by activating another policy set. Because of this copying process, you should not assume that the definitions in the policy set that was last activated are the same as those in the ACTIVE policy set; the original policy set may have been modified since activation.

Privilege Class

Any administrator can issue this command.

Syntax



Parameters

domainname

Specifies the name of the policy domain to be queried. This parameter is optional. You can use a pattern-matching expression to specify this name. All matching policy domains are displayed. If you do not specify a value for this parameter, all policy domains are displayed.

Format=*formatvalue*

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed for the specified policy domain.

Detailed

Specifies that complete information is displayed for the specified policy domain.

Examples

Task 1

Display partial information for all policy domains.

QUERY DOMAIN

Command: Query D0main

Output:

Policy Domain Name	Activated Policy Set	Activated Default Mgmt Class	Number of Registered Nodes	Description
DOMAIN3	PGSET	MCLASS1	1	Department 1
WORLD1	TOM1	TOM1	8	Department 2

See Task 2 for parameter definitions.

Task 2

Display information in full detail for the EMPLOYEE_RECORDS policy domain.

Command: Query D0main employee_records Format=Detailed

Output:

```
Policy Domain Name: EMPLOYEE_RECORDS
Activated Policy Set: VACATION_
Activation Date/Time: 04/13/1993 09:11:10
Days Since Activation: <1
Activated Default Mgmt Class: ACTIVEFILES
Number of Registered Nodes: 0
Description:
Backup Retention (Grace Period): 30
Archive Retention (Grace Period): 60
Last Update by (administrator): PICASSO
Last Update Date/Time: 04/13/1993 09:11:10
```

where:

Policy Domain Name

Specifies the name of the policy domain.

Activated Policy Set

Specifies the name of the policy set that was last activated in the domain.

Activation Date/Time

Specifies the date and time that the policy set was activated.

Days Since Activation

Specifies the number of days since the policy set was activated.

Activated Default Mgmt Class

Specifies the assigned default management class for the policy set.

Number of Registered Nodes

Specifies the number of client nodes registered to the policy domain.

Description

Specifies a description of the policy domain.

Backup Retention (Grace Period)

Specifies the number of days to retain backup versions that are no longer on the client's system. The backup retention value is used by ADSM to manage inactive versions of files when any of the following conditions are present:

- A file's management class is changed, but neither the new management class nor default management class contains a backup copy group.
- The management class to which a file is bound no longer exists, and the default management class does not contain a backup copy group.
- The backup copy group is deleted from the management class to which a file is bound and the default management class does not contain a backup copy group.

Archive Retention (Grace Period)

Specifies the number of days to retain an archive file that meets either of the following conditions:

- The management class to which a file is bound no longer exists, and the default management class does not contain an archive copy group.
- The archive copy group is deleted from the management class to which a file is bound and the default management class does not contain an archive copy group.

Last Update by (administrator)

Specifies the name of the administrator that defined or most recently updated the policy domain.

Last Update Date/Time

Specifies when the administrator defined or most recently updated the storage pool.

Related Commands

Table 90 provides references to related commands.

Table 90. Commands Related to QUERY DOMAIN

Command	Description	Reference
COPY DOMAIN	Creates a copy of an existing policy domain.	Page 70
DEFINE DOMAIN	Defines a policy domain that can be assigned to one or more clients.	Page 107
DELETE DOMAIN	Deletes a policy domain along with any policy objects in the policy domain.	Page 162
UPDATE DOMAIN	Changes the attributes of a policy domain.	Page 458

QUERY EVENT (Query Scheduled and Completed Events)

Use the QUERY EVENT command to display scheduled and completed events. The QUERY EVENT command takes two forms, depending on whether the query applies to scheduled client operations or scheduled administrative commands. The syntax and parameters for each form are defined separately.

Querying Events for Client Schedules

Use the QUERY EVENT command to display scheduled and completed events for selected clients. You can specify a time range to limit output to events whose scheduled start times fall within the range.

Each scheduled client operation is called an event. The server tracks each scheduled event and records the results of each function in the database. An output record is created whenever processing of a scheduled command is started or missed.

You can request one of two different types of queries: *general* reporting or *exception* reporting. To request a general report of scheduled events, specify this command with EXCEPTIONSONLY=NO (the default). For each schedule that matches the specified policy domain and schedule name, the expected startup windows are determined within the specified time range. ADSM provides output data that indicates scheduled start times, along with any actual execution information which can be obtained from existing output records.

To request an exception report, specify EXCEPTIONSONLY=YES. ADSM provides information about scheduled operations that have *not* executed successfully, either because they have not executed as scheduled or because they have failed during execution. This mode differs from EXCEPTIONSONLY=NO, in that output is restricted as follows:

- The only events that are displayed are those that meet *one* of the following conditions:
 - The scheduled startup window has elapsed and the schedule has not begun execution.
 - The client has notified the server that the scheduled command has failed during execution and all allowed retry attempts have been exhausted.
 - The schedule began, but has not completed execution and the schedule period has elapsed since the scheduled start time.
- Information is only displayed for events that should have an output record. Thus, output records are not displayed if the records are missing but are eligible for deletion through *either* of the following:
 - Removal of output records whose retention period has elapsed. For more information on setting the event retention period, see “SET EVENTRETENTION (Specify the Retention Period for Event Records)” on page 412.
 - Execution of a DELETE EVENT command.

No information is displayed for an event unless both of the following conditions are met:

- The last modification date for the schedule occurred before the end of the startup window for the event.
- The date on which the client was associated with the schedule occurred before the end of the startup window for the event.

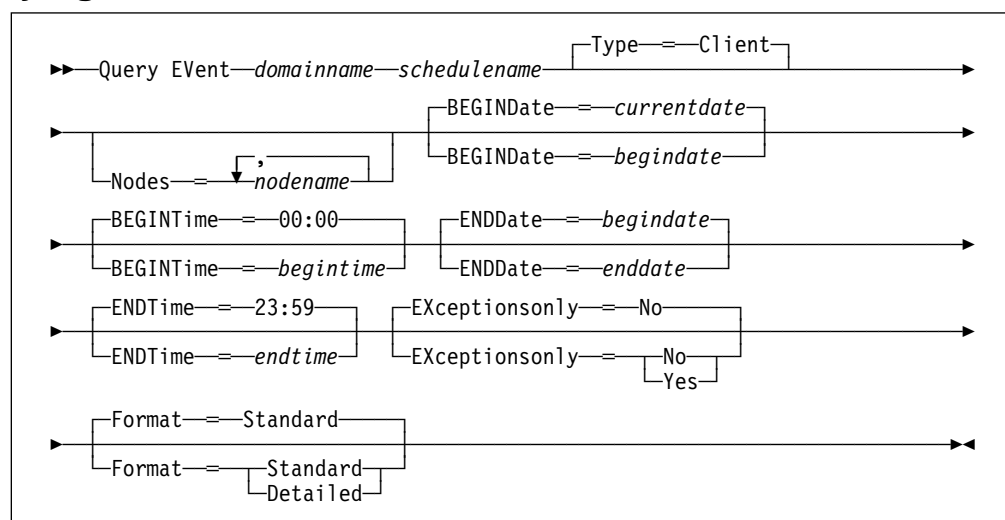
Processing of this command can require intensive computations, especially if multiple startup windows must be calculated. To minimize the processing time when querying events:

- Minimize the length of the time range for querying events.
- Restrict the policy domains, schedule names, and client node names to those for which information is required.

Privilege Class

Any administrator can issue this command.

Syntax for Querying Events for Client Schedules



Parameters for Querying Events for Client Schedules

domainname

Specifies the name of the policy domain to which the schedules belong. This parameter is required. You can use a pattern-matching expression to specify this name. All matching policy domains are queried.

schedulename

Specifies the name of the schedule for which events are displayed. This parameter is required. You can use a pattern-matching expression to specify this name. All matching schedule names are queried.

Type=Client

Specifies that this command should query events for client schedules. This parameter is optional. The default value is CLIENT. Administrative command schedule events are queried by specifying TYPE=ADMINISTRATIVE.

Nodes=nodename

Specifies the name of the client node that belongs to the specified policy domain for which events are displayed. You can specify more than one client name by separating each name with a comma, with no intervening spaces.

QUERY EVENT

You can use a pattern-matching expression to specify a name. All matching clients are queried. If you do not specify a client name, events for all clients matching the DOMAINNAME and SCHEDULENAME parameters are displayed.

BEGINDate=*date*

Specifies the beginning date of the time range for the events to be displayed. All events that were scheduled to start during this time are displayed. This parameter is optional. The default value is the current date.

Possible values are:

MM/DD/YYYY

Specifies the date for the beginning of the time range in MM/DD/YYYY format (for example, 01/31/1995).

TODAY

Specifies that the beginning of the time range is the current date.

TODAY-*numdays* or -*numdays*

Specifies that the beginning of the time range is a specific number of days before the current date. The maximum value for the number of days is 9999.

For example, to query events scheduled to start during the past 7 days, you can specify one of the following parameter combinations:

BEGINDATE=TODAY-7 ENDDATE=TODAY

or

BEGINDATE=-7 ENDDATE=TODAY

TODAY+*numdays* or +*numdays*

Specifies that the beginning of the time range is a specific number of days after the current date. The maximum value for the number of days is 9999.

For example, to query events scheduled to start two days from now, you can specify one of the following parameters:

BEGINDATE=TODAY+2

or

BEGINDATE=+2

BEGINTime=*time*

Specifies the beginning time of the range for events to be displayed, in HH:MM format. All events scheduled to start during this time are displayed. This parameter is optional. The default value is 00:00.

ENDDate=*date*

Specifies the ending date of the time range for events to be displayed. All events that were scheduled to start during this time are displayed. This parameter is optional. The default value is the value used for the BEGINDATE parameter.

Possible values are:

MM/DD/YYYY

Specifies the date for the end of the time range in MM/DD/YYYY format (for example, 01/31/1995).

TODAY

Specifies that the end of the time range is the current date.

TODAY-*numdays* or -*numdays*

Specifies that the end of the time range is a specific number of days before the current date. The maximum value for the number of days is 9999.

For example, to query events scheduled to start during a one-week period that ended yesterday, you can specify one of the following parameter combinations:

BEGINDATE=TODAY-8 ENDDATE=TODAY-1
or
BEGINDATE=-8 ENDDATE=-1

TODAY+*numdays* or +*numdays*

Specifies that the end of the time range is a specific number of days after the current date. The maximum value for the number of days is 9999.

For example, to query events scheduled to start during the next week, you can specify one of the following parameter combinations:

BEGINDATE=TODAY ENDDATE=TODAY+7
or
BEGINDATE=TODAY ENDDATE=+7

ENDTime=*time*

Specifies the ending time of the range for events to be displayed, in HH:MM format. All events scheduled to start during this time are displayed. This parameter is optional. The default value is 23:59.

EXceptiononly=*exceptionvalue*

Specifies the type of query information you want on scheduled or completed events. This parameter is optional. The default value is NO. Possible values are:

No

Specifies that information on past and projected events is displayed.

Yes

Specifies that only those events which failed or did not execute as scheduled will be displayed.

Format=*formatvalue*

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information for events is displayed.

Detailed

Specifies that complete information for events is displayed.

Querying Events for Administrative Command Schedules

Use the QUERY EVENT command to display scheduled and completed events for selected administrative command schedules.

Each scheduled administrative command is called an event. The server tracks each scheduled event and records the results of each function in the database. An output record is created whenever processing of a scheduled command is started or missed.

You can request one of two different types of queries: *general* reporting or *exception* reporting. To request a general report of scheduled events, specify this command with EXCEPTIONSONLY=NO (the default). For each schedule that matches the specified schedule name, the expected startup windows are determined within the specified time range. ADSM provides output data that indicates scheduled start times, along with any actual execution information which can be obtained from existing output records.

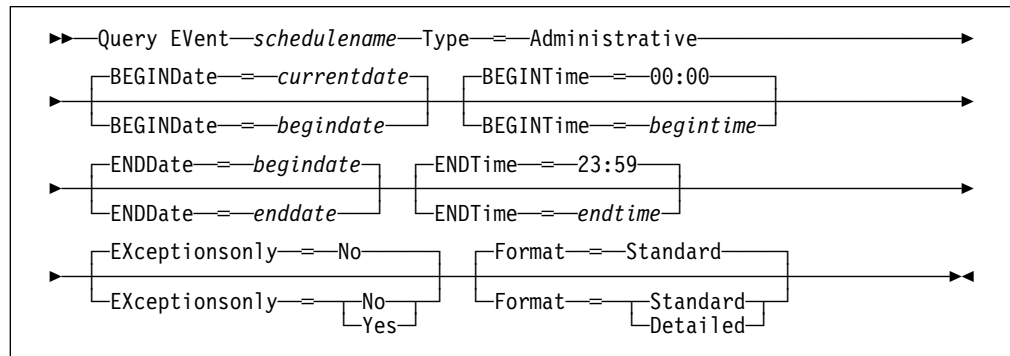
Processing of this command can require intensive computations, especially if multiple startup windows must be calculated. To minimize the processing time when querying events:

- Minimize the length of the time range for querying events.
- Restrict the schedule names to those for which information is required.

Privilege Class

Any administrator can issue this command.

Syntax for Querying Events for Administrative Command Schedules



Parameters for Querying Events for Administrative Command Schedules

schedulename

Specifies the name of the schedule for which events are displayed. This parameter is required. You can use a pattern-matching expression to specify this name. All matching schedule names are queried.

Type=**Administrative**

Specifies that you want to query events for administrative command schedules. This parameter is required for querying administrative command schedules.

BEGINDate=*date*

Specifies the beginning date of the time range for the events to be displayed. All events that were scheduled to start during this time are displayed. This parameter is optional. The default value is the current date.

Possible values are:

MM/DD/YYYY

Specifies the date for the beginning of the time range in MM/DD/YYYY format (for example, 01/31/1995).

TODAY

Specifies that the beginning of the time range is the current date.

TODAY-*numdays* **or** **-***numdays*

Specifies that the beginning of the time range is a specific number of days before the current date. The maximum value for the number of days is 9999.

For example, to query events scheduled to start during the past 7 days, you can specify one of the following parameter combinations:

BEGINDATE=TODAY-7 ENDDATE=TODAY
 or
 BEGINDATE=-7 ENDDATE=TODAY

TODAY+*numdays* **or** **+***numdays*

Specifies that the beginning of the time range is a specific number of days after the current date. The maximum value for the number of days is 9999.

For example, to query events scheduled to start two days from now, you can specify one of the following parameters:

BEGINDATE=TODAY+2
 or
 BEGINDATE=+2

BEGINTime=*time*

Specifies the beginning time of the range for events to be displayed, in HH:MM format. All events scheduled to start during this time are displayed. This parameter is optional. The default value is 00:00.

ENDDate=*date*

Specifies the ending date of the time range for events to be displayed. in MM/DD/YYYY format. All events that were scheduled to start during this time are displayed. This parameter is optional. The default value is the value used for the BEGINDATE parameter.

Possible values are:

MM/DD/YYYY

Specifies the date for the end of the time range in MM/DD/YYYY format (for example, 01/31/1995).

TODAY

Specifies that the end of the time range is the current date.

QUERY EVENT

TODAY*-numdays or -numdays*

Specifies that the end of the time range is a specific number of days before the current date. The maximum value for the number of days is 9999.

For example, to query events scheduled to start during a one-week period that ended yesterday, you can specify one of the following parameter combinations:

BEGINDATE=TODAY-8 ENDDATE=TODAY-1
or
BEGINDATE=-8 ENDDATE=-1

TODAY*+numdays or +numdays*

Specifies that the end of the time range is a specific number of days after the current date. The maximum value for the number of days is 9999.

For example, to query events scheduled to start during the next week, you can specify one of the following parameter combinations:

BEGINDATE=TODAY ENDDATE=TODAY+7
or
BEGINDATE=TODAY ENDDATE=+7

ENDTime*=time*

Specifies the ending time of the range for events to be displayed, in HH:MM format. All events scheduled to start during this time are displayed. This parameter is optional. The default value is 23:59.

EXceptiononly*=exceptionvalue*

Specifies the type of query information you want on scheduled or completed events. This parameter is optional. The default value is NO. Possible values are:

No

Specifies that information on past and projected events is displayed.

Yes

Specifies that only those events which failed or did not execute as scheduled will be displayed.

Format*=formatvalue*

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information for events is displayed.

Detailed

Specifies that complete information for events is displayed.

Examples

Task 1

Display partial information for all events scheduled for an administrative schedule named DOSADMIN. Specify administrative command schedules only. Limit the query to events scheduled for March 30, 1995 (03/30/1995).

Command: Query Evt * Type=Administrative BEGINDate=03/30/1995
ENDDate=03/30/1995

Output:

Scheduled Start	Actual Start	Schedule Name	Status
-----	-----	-----	-----
03/30/1995 00:00:00	03/30/1995 00:00:01	DOSADMIN	Completed
03/30/1995 04:00:00	03/30/1995 04:00:01	DOSADMIN	Completed
03/30/1995 08:00:00	03/30/1995 08:00:01	DOSADMIN	Completed
03/30/1995 12:00:00		DOSADMIN	Future
03/30/1995 16:00:00		DOSADMIN	Future
03/30/1995 20:00:00		DOSADMIN	Future

See Task 3 for parameter definitions.

Task 2

Display partial information for all events scheduled for DOMAIN1 that did not execute successfully. Limit the search to the client named JOE. Limit the events displayed to those scheduled to occur from July 11, 1994 (07/11/1994) to July 12, 1994 (07/12/1994).

Command: Query Evt domain1 * Nodes=joe BEGINDate=07/11/1994
ENDDate=07/12/1994 EXceptiononly=Yes

Output:

Scheduled Start	Actual Start	Schedule Name	Node Name	Status
-----	-----	-----	-----	-----
07/11/1994 01:00:00	07/11/1994 01:13:55	BACK1	JOE	Completed
07/12/1994 01:00:00		DAILYBKP	JOE	Missed

See Task 3 for parameter definitions.

Task 3

Display complete information on all events that are scheduled for execution by client JOE in DOMAIN1 any time from July 11, 1994 (7/11/94) through July 12, 1994 (7/12/94).

Command: Query Evt domain1 * Nodes=joe BEGINDate=07/11/1994
ENDDate=07/12/1994 Format=Detailed

Output:

```
Policy Domain Name: DOMAIN1
  Schedule Name: DAILY_BACKUP
    Node Name: JOE
      Scheduled Start: 07/11/1994 01:00:00
        Actual Start:
          Completed:
            Status: Missed
              Result: 0

Policy Domain Name: DOMAIN1
  Schedule Name: DAILY_BACKUP
    Node Name: JOE
      Scheduled Start: 07/11/1994 01:00:00
        Actual Start: 07/11/1994 01:13:55
          Completed: 07/11/1994 01:18:27
            Status: Completed
              Result: 0
```

where:

Policy Domain Name

Specifies the name of the policy domain to which the schedule is assigned.

Schedule Name

Specifies the name of the schedule that has initiated this event.

Node Name

Specifies the client scheduled to perform the operation.

Scheduled Start

Specifies the scheduled starting date and time for the event. The scheduled start time displayed does not include any randomization that can occur through the use of the SET RANDOMIZE command.

Actual Start

Specifies the actual date and time at which the client began execution of the scheduled operation. No information is displayed if the schedule has not started executing.

Completed

Specifies the date and time the scheduled event is completed.

Status

Specifies the completion status of the event. Possible values are:

Completed

Specifies that the scheduled event has been completed.

Missed

Specifies that the scheduled startup window for this event has passed and the schedule has not begun.

Failed

Specifies that the client reported a failure in executing the scheduled operation and successive retries have not succeeded.

Started

Specifies that processing began for the event.

Restarted

Specifies that the client has retried execution of the scheduled operation.

Severed

Specifies that communication with the client was severed before the event could complete.

Future

Specifies that the beginning of the startup window for the event is in the future.

Pending

Specifies that the QUERY EVENT command was issued during the startup window for the event, but execution of the scheduled operation has not yet begun.

Uncertain

Specifies that the state of the event cannot be determined, for example, whether the event record has been deleted.

Result

Specifies the return code from the client that identifies whether the schedule has executed successfully. Refer to the appropriate user's guide for more information on the value for this return code.

Related Commands

Table 91 provides references to related commands.

<i>Table 91. Commands Related to QUERY EVENT</i>		
Command	Description	Reference
DELETE EVENT	Deletes event records prior to a specified date and time.	Page 164
SET EVENTRETENTION	Specifies the number of days to retain event records.	Page 412
SET RANDOMIZE	Specifies the randomization of start times within a window for schedules in client-polling mode.	Page 426

QUERY FILESPACE (Query One or More File Spaces)

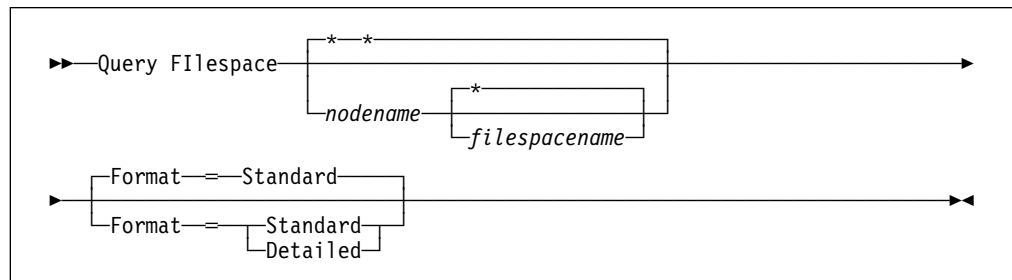
Use the QUERY FILESPACE command to display information about one or more file spaces that belong to a client node. The information displayed via the QUERY FILESPACE command reflects data as of the last incremental backup.

File space names are case sensitive. You can use the QUERY FILESPACE command to determine the correct capitalization for the file space to be queried.

Privilege Class

Any administrator can issue this command.

Syntax



Parameters

nodename

Specifies the name of the client node to which the file space belongs. This parameter is optional. You can use a pattern-matching expression to specify this name. All matching client node names are included. If you do not specify a value for this parameter, all client nodes are queried.

You must specify a value for this parameter if you specify FILESPACE NAME.

filespace name

Specifies the name of the file space to be queried. This parameter is optional. You can use a pattern-matching expression to specify this name. All matching file spaces are displayed. If you do not specify a value for this parameter, all file spaces are queried.

Format=*formatvalue*

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed for the specified file space.

Detailed

Specifies that complete information is displayed for the specified file space.

Examples

Task 1

Query all file spaces associated with all client nodes.

Command: Query Filespace

Output:

Node Name	Filespace Name	Platform	Filespace Type	Capacity (MB)	%Util
GOODPELL	OS2	OS/2	FAT	0.0	0.0
HT	C_DISK	OS/2	HPFS	0.0	0.0
TANG	C_DISK	OS/2	FAT	200.8	80.8
TANG	DRIVE_D	OS/2	FAT	99.8	79.4
VIKINGS	/home/cpark	AIX	JFS	0.0	0.0

where:

Node Name

Specifies the name of the client node.

Filespace Name

Specifies the name of the file space for the client node.

Platform

Specifies the platform for the client node.

Filespace Type

Specifies the type of file space. Possible values are AFS, DOS, FAT, HPFS, JFS.

Capacity (MB)

Specifies the amount of space assigned to this file space, in megabytes, on the client node. This value applies to the actual file systems and media on the client node's machine.

%Util

Specifies, as a percentage of the capacity of the file space, the amount of occupied space in the file space, on the client node. This value applies to the actual file systems and media on the client node's machine.

Task 2

Display detailed information on the OS2 file space that belongs to the client node SHUNA.

Command: Query Filespace shuna OS2 Format=Detailed

QUERY FILESPACE

Output:

```
Node Name: SHUNA
Filespace Name: OS2
Platform: OS/2
Filespace Type:
Capacity (MB): 0.0
%Util: 0.0
Last Backup Start Date/Time:
Days Since Last Backup Started:
Last Backup Completion Date/Time:
Days Since Last Backup Completed:
```

where:

Node Name

Specifies the name of the client node.

Filespace Name

Specifies the name of the file space for the client node.

Platform

Specifies the platform for the client node.

Filespace Type

Specifies the type of file space. Possible values are AFS, DOS, FAT, HPFS, JFS.

Capacity (MB)

Specifies the amount of space assigned to this file space, in megabytes, on the client node.

%Util

Specifies, as a percentage of the capacity for the file space, the amount of space that is occupied in the file space on the client node.

Last Backup Start Date/Time

Specifies the date and time that the last incremental backup of the file space was initiated.

Days Since Last Backup Started

Specifies the number of days since the last incremental backup of the file space was initiated.

Last Backup Completion Date/Time

Specifies the completion date and time of the last incremental backup of the file space.

Days Since Last Backup Completed

Specifies the number of days since the last incremental backup of the file space completed.

Related Commands

Table 92 provides references to related commands.

Table 92. Commands Related to QUERY FILESPACE

Command	Description	Reference
DELETE FILESPACE	Deletes data associated with one or more client's filespaces.	Page 167
REGISTER NODE	Defines a client to the server and sets options for that user.	Page 374
REMOVE NODE	Removes a client from the list of registered nodes for a specific policy domain.	Page 378
UPDATE NODE	Changes the attributes associated with a client node.	Page 464

QUERY LICENSE (Display License Information)

Use the QUERY LICENSE command to display information about the results of the latest license audit and the server license terms. The information displayed indicates whether the server is in compliance with the license terms. This information is useful for determining how to manage license terms or policy changes.

Privilege Class

Any administrator can issue this command.

Syntax

```
▶▶—Query LICense—————▶▶
```

Parameters

None.

Examples

Task

Display information about license terms and compliance information.

Command: Query LICense

Output:

```
                Last License Audit: 05/08/1995 16:37:29
                Registered Client Nodes: 20
                Licensed Client Nodes: 50
                Are UNIX clients registered ?: NO
                Are UNIX clients licensed ?: NO
                Are DeskTop clients registered ?: YES
                Are DeskTop clients licensed ?: YES
                Are OEMVS clients registered ?: YES
                Are OEMVS clients licensed ?: YES
                Is space management in use on the server ?: YES
                Is space management licensed on the server ?: YES
                Device support module 1 licensed ?: NO
                Device support module 2 licensed ?: NO
                Device support module 3 licensed ?: YES
                Device support module 4 licensed ?: NO
                Secondary Server Attachment licensed ?: NO
                Server License Compliance: Valid
```

where:

Last License Audit On

Specifies the date and time when the last license audit took place. The displayed information is valid as of this date and time.

Registered Client Nodes

Specifies the total number of registered client nodes.

Licensed Client Nodes

Specifies the total number of client nodes that can be registered as defined in the license terms.

Are UNIX clients registered?

Specifies whether UNIX clients are registered with the server.

Are UNIX clients licensed?

Specifies whether the server is licensed to support UNIX clients.

Are DeskTop clients registered?

Specifies whether desktop clients (such as DOS) are registered with the server.

Are DeskTop clients licensed?

Specifies whether the server is licensed to support desktop clients.

Are OEMVS clients registered?

Specifies whether OpenEdition MVS clients are registered with the server.

Are OEMVS clients licensed?

Specifies whether the server is licensed to support OpenEdition MVS clients.

Is space management in use on the server?

Specifies whether HSM clients are attached to the server.

Is space management licensed on the server?

Specifies whether the server is licensed to support HSM clients.

Device support module *n* licensed?

Specifies whether the server is licensed for the device support module.

Secondary Server Attachment licensed?

Specifies whether the server is licensed for support of devices through the secondary server attachment feature.

Server License Compliance:

Specifies whether the ADSM server is in compliance, pending failure, or out of compliance with registered client node license terms. The server is in compliance when the registered client nodes are less than 100 percent of the license terms. When registered client nodes are between 100 percent and 120 percent of the license terms, a pending failure condition is present. When the registered client nodes exceed 120 percent of the license terms, the server is out of compliance.

Related Commands

Table 93 provides references to related commands.

Table 93. Commands Related to QUERY LICENSE

Command	Description	Reference
AUDIT LICENSES	Checks for compliance with defined licenses.	Page 43
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.	Page 342
REGISTER LICENSE	Registers a new license with the ADSM server.	Page 372
SET LICENSEAUDITPERIOD	Specifies the number of days between automatic license audits.	Page 414


QUERY LOG (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.

Privilege Class

Any administrator can issue this command.

Syntax



Parameters

Format=*formatvalue*

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed for the recovery log.

Detailed

Specifies that complete information is displayed for the recovery log.

Examples

Task 1

Display information about the recovery log.

Command: Query LOG

Output:

Available Space (MB)	Assigned Capacity (MB)	Maximum Extension (MB)	Maximum Reduction (MB)	Page Size (bytes)	Total Usable Pages	Used Pages	%Util	Max. %Util
16	16	0	12	4,096	4,096	247	6.0	6.0

See Task 2 for parameter definitions.

Task 2

Display information about the recovery log in complete detail.

Command: Query LOG Format=Detailed

Output:

```
Available Space (MB): 12
Assigned Capacity (MB): 12
Maximum Extension (MB): 0
Maximum Reduction (MB): 0
  Page Size (bytes): 4,096
  Total Usable Pages: 2,560
    Used Pages: 243
    %Util: 9.5
    Max. %Util: 9.9
  Physical Volumes: 1
  Log Pool Pages: 24
  Log Pool Pct. Util: 2.94
  Log Pool Pct. Wait: 0.00
Cumulative Consumption (MB): 57.00
Consumption Reset Date/Time: 02/20/1995 13:23:47
```

where:

Available Space (MB)

Specifies the amount of space available in the recovery log, in megabytes.

Assigned Capacity (MB)

Specifies the assigned storage capacity for the recovery log, in megabytes.

Maximum Extension (MB)

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

Maximum Reduction (MB)

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

Page Size (bytes)

Specifies the page size, in bytes.

Total Usable Pages

Specifies the total number of pages available for transactions in the recovery log. A page is a unit of allocation within recovery log volumes.

Used Pages

Specifies the total number of pages currently in use by the recovery log.

%Util

Specifies the current utilization as a percentage of the assigned capacity.

Max. %Util

Specifies the maximum occupancy for the recovery log, as a percentage of the available space. This value specifies the highest capacity occupied by the recovery log since the last reset of this statistic. You can reset this value with the RESET LOGMAXUTILIZATION command.

Physical Volumes

Specifies the total number of recovery log volumes defined to the recovery log, including volume copies.

Log Pool Pages

Specifies the number of pages in the recovery log buffer pool. This value is determined by the value for the LOGPOOLSIZE option. The default value for the LOGPOOLSIZE option is 128 kilobytes (32 pages).

Log Pool Pct. Util

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

Log Pool Pct. Wait

Specifies, as a percentage of all recovery log buffer pool requests, the number of requests for a recovery log buffer pool page that is not available (because all buffer pool pages are waiting to write to the recovery log). If this number is greater than 0.0, you should increase the size of the buffer pool with the LOGPOOLSIZE option.

Cumulative Consumption

Specifies the amount of log space used by the server since the last reset of this statistic. You can reset this statistic with the RESET LOGCONSUMPTION command.

Consumption Reset Date/Time

Specifies the date and time of the last reset of the cumulative consumption statistic for the recovery log. You can reset this statistic with the RESET LOGCONSUMPTION command.

Related Commands

Table 94 provides references to related commands.

Table 94 (Page 1 of 2). Commands Related to QUERY LOG

Command	Description	Reference
DEFINE LOGCOPY	Creates a volume copy of a recovery log volume.	Page 110
DEFINE LOGVOLUME	Defines a disk volume to the recovery log.	Page 112
EXTEND DB	Extends usable database storage within the volumes previously allocated for its use.	Page 213

QUERY LOG

Table 94 (Page 2 of 2). Commands Related to QUERY LOG

Command	Description	Reference
EXTEND LOG	Extends usable log storage within the volumes previously allocated for its use.	Page 215
REDUCE DB	Attempts to free up a database volume by reducing usable storage within the volumes defined.	Page 366
REDUCE LOG	Attempts to free up a log volume by reducing usable storage within the volumes defined.	Page 368
RESET LOGCONSUMPTION	Resets the statistic on recovery log space used.	Page 389
RESET LOGMAXUTILIZATION	Resets the maximum utilization statistic for the recovery log.	Page 390


QUERY LOGVOLUME (Display Information on One or More 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.

Privilege Class

Any administrator can issue this command.

Syntax



Parameters

volname

Specifies the name of the recovery log volume to be queried. This parameter is optional. You can use a pattern-matching expression to specify this name. All matching recovery log volumes are displayed. If you do not specify a volume name, all volumes are displayed.

Format=*formatvalue*

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed for the specified recovery log volume.

Detailed

Specifies that complete information is displayed for the specified recovery log volume.

Examples

Task 1

Display information about all recovery log volumes.

Command: Query LOGVolume

Output:

Volume Name (Copy 1)	Copy Status	Volume Name (Copy 2)	Copy Status	Volume Name (Copy 3)	Copy Status
0301	Sync'd		Undef- ined		Undef- ined
0302	Sync'd		Undef- ined		Undef- ined

QUERY LOGVOLUME

See Task 2 for parameter definitions.

Task 2

Display detailed information about the recovery log volume named HARVEY.ADSM.PINV.

Command: Query LOGVolume harvey.adsm.pinv Format=Detailed

Output:

```
Volume Name (Copy 1): HARVEY.ADSM.PINV
    Copy Status: Sync'd
Volume Name (Copy 2): GLEDDIE.ADSM.PINV
    Copy Status: Sync'd
Volume Name (Copy 3):
    Copy Status: Undefined
Available Space (MB): 36
Allocated Space (MB): 36
    Free Space (MB): 0
```

where:

Volume Name (Copy *n*)

Specifies the name of the database volume. The output displays status of the three volume copies (Copy 1, Copy 2, and Copy 3).

Copy Status

Specifies the status of the database volume. Possible values are:

Offline

The volume is varied offline.

Stale

The volume is not synchronized with the other volumes in the recovery log and cannot be used for storage.

Sync'd

The volume is synchronized with the other volumes in the recovery log and can be used for storage.

Undefined

No volume has been defined.

Available Space (MB)

Specifies the total amount of space that is available, both allocated and free, for the recovery log, in megabytes.

Allocated Space

Specifies the amount of space allocated for the recovery log, in megabytes.

Free Space

Specifies the amount of unassigned space in the recovery log that can be allocated for use, in megabytes.

Related Commands

Table 95 provides references to related commands.

Table 95. Commands Related to QUERY LOGVOLUME

Command	Description	Reference
DEFINE LOGCOPY	Creates a volume copy of a recovery log volume.	Page 110
DEFINE LOGVOLUME	Defines a disk volume to the recovery log.	Page 112
DELETE LOGVOLUME	Deletes a volume from the recovery log.	Page 170
VARY	Specifies whether a disk volume is available to the server for use.	Page 504

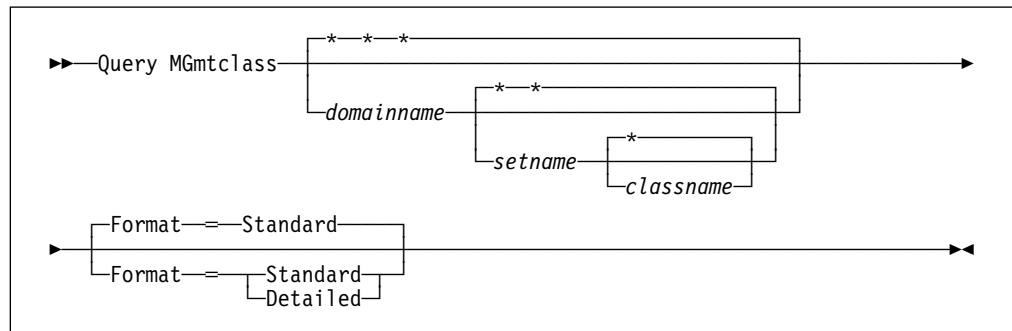
QUERY MGMTCLASS (Query a Management Class)

Use the QUERY MGMTCLASS command to display information on an existing management class. This command is useful for determining which management classes are available, including the default management class, for a policy domain. To identify the default management class, use the QUERY DOMAIN command.

Privilege Class

Any administrator can issue this command.

Syntax



Parameters

domainname

Specifies the name of the policy domain associated with the management class to be queried. This parameter is optional. You can use a pattern-matching expression to specify this name. All matching policy domains are queried. If you do not specify a value for this parameter, all policy domains are queried. You must specify this parameter when querying an explicitly named management class.

setname

Specifies the name of the policy set associated with the management class to be queried. This parameter is optional. You can use a pattern-matching expression to specify this name. All matching policy sets are queried. If you do not specify a value for this parameter, all policy sets are queried. You must specify this parameter when querying an explicitly named management class.

classname

Specifies the name of the management class associated with the policy set to be queried. This parameter is optional. You can use a pattern-matching expression to specify this name. All matching management classes are queried. If you do not specify a value for this parameter, all management classes are queried. You must specify this parameter when querying an explicitly named management class.

Format=*formatvalue*

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed for the specified management class.

Detailed

Specifies that complete information is displayed for the specified management class.

Examples

Task 1

Query all management classes for all policy domains. Create the output in standard format.

Command: Query Mgmtclass

Output:

Policy Domain Name	Policy Set Name	Mgmt Class Name	Default Mgmt Class ?	Description
EMPLOYEE-RECORDS	ACTIVE	ACTIVEFILES	Yes	
EMPLOYEE-RECORDS	VACATION	ACTIVEFILES	Yes	
PROG1	SUMMER	DEFAULT1	Yes	Technical Support Management Class
PROG1	SUMMER	MCLASS1	No	Technical Support Management Class
STANDARD	ACTIVE	STANDARD	Yes	Installed default management class
STANDARD	STANDARD	STANDARD	Yes	Installed default management class

See Task 2 for parameter definitions. **Task 2**

Query the ACTIVEFILES management class that is assigned to the VACATION policy set of the EMPLOYEE_RECORDS policy domain. Create the output in detailed format.

Command: Query Mgmtclass employee_records vacation activefiles
Format=Detailed

Output:

```
Policy Domain Name: EMPLOYEE_RECORDS
Policy Set Name: VACATION
Mgmt Class Name: ACTIVEFILES
Default Mgmt Class?: Yes
Description:
Space Management Technique: Automatic
Auto-Migrate on Non-Use: 0
Migration Requires Backup?: Yes
Migration Destination: SPACEMGPOOL
Last Update by (administrator): SERVER_CONSOLE
Last Update Date/Time: 03/29/1995 15:00:01
```

where:

Policy Domain Name

Specifies the name of the policy domain.

Policy Set Name

Specifies the name of the activated policy set for the policy domain.

Mgmt Class Name?

Specifies the name of the management class.

Default Mgmt Class

Specifies whether the management class is defined as the default management class for the policy set.

Description

Specifies a description of the management class.

Space Management Technique

Specifies the space management technique for the management class.

Auto-Migrate on Non-Use

Specifies the number of days that must elapse since the file was last accessed before it is eligible for automatic migration.

Migration Requires Backup?

Specifies whether a backup version of a file must exist before a file can be migrated.

Migration Destination

Specifies the name of an ADSM storage pool that is a destination for migrated files.

Last Update by (administrator)

Specifies the name of the administrator that defined or most recently updated the management class.

Last Update Date/Time

Specifies the date and time that the administrator defined or most recently updated the management class.

Related Commands

Table 96 provides references to related commands.

Table 96. Commands Related to QUERY MGMTCLASS

Command	Description	Reference
COPY MGMTCLASS	Creates a copy of an existing management class.	Page 73
DEFINE MGMTCLASS	Defines a management class within the policy domain and policy set specified.	Page 114
DELETE MGMTCLASS	Deletes a management class and its copy groups from a policy domain and policy set.	Page 172
QUERY DOMAIN	Displays information about one or more policy domains.	Page 287
UPDATE MGMTCLASS	Changes the attributes of a management class.	Page 460


QUERY MOUNT (Display Information on Mounted Sequential Access Volumes)

Use the QUERY MOUNT command from the administrative client to request information about the status of one or all sequential access volumes that are mounted.

Privilege Class

Any administrator can issue this command.

Syntax



Parameters

volname

Specifies the name of the mounted sequential access volume that you want to query. This parameter is optional. You can use a pattern-matching expression to specify this name. All matching mounted volumes are queried. If you do not specify a sequential access volume name, information on all volumes is displayed.

Examples

Task 1

Display information on the mounted sequential access media volumes named BTV005.

Command: Query M0unt btv005

Output:

```

ANR52891I CARTRIDGE volume BTV005 mounted R/W on 0FC6,
status: "In Use."
  
```

Task 2

Display information on all mounted sequential access media volumes with names that begin with "H."

Command: Query M0unt h*

Output:

```

ANR52890I CARTRIDGE volume HTV007 mounted R/W on 0FC6,
status: "Idle."
ANR5290I CARTRIDGE volume HTV003 mounted R/W on 0FC8,
status: "In Use."
  
```

Related Commands

Table 97 provides references to related commands.

Table 97. Commands Related to QUERY MOUNT

Command	Description	Reference
DISMOUNT VOLUME	Dismounts a sequential, removable volume by the volume name.	Page 191

QUERY NODE (Query One or More Client Nodes)

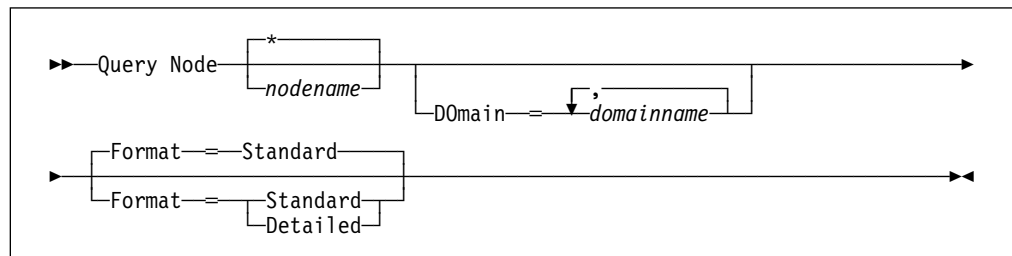
Use the QUERY NODE command to display information about one or more registered client nodes. This command provides client node information such as:

- The name of the client node
- Date of last access to the server
- Date password was last changed
- The client node's policy domain
- The client node's address and communication protocol
- The communication method used by the client node

Privilege Class

Any administrator can issue this command.

Syntax



Parameters

nodename

Specifies the name of the client node to be queried. You can use a pattern-matching expression to specify this name. All matching client nodes are queried. If you do not specify a value for this parameter, all client nodes are queried.

Domain=domainname

Specifies a list of policy domains that limit the client node query. Only nodes that are assigned to one of the specified policy domains are displayed. This parameter is optional. The items in the list are separated by commas, with no intervening spaces. You can use a pattern-matching expression to specify a domain. All clients assigned to a matching domain will be displayed. If you do not specify a value for this parameter, all policy domains are included in the query.

Format=formatvalue

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed for the specified client nodes.

Detailed

Specifies that complete information is displayed for the specified client nodes.

Examples

Task 1

Display information on all registered client nodes.

Command: Query Node

Output:

Node Name	Platform	Policy Domain Name	Days Since Last Access	Days Since Password Set	Locked?
GOODELL	(?)	STANDARD	<1	<1	No
HT	OS/2	STANDARD	2	8	No
RCHANG	AIX	STANDARD	10	10	No
SHUNA	OS/2	STANDARD	7	7	No
SHELLBEE	AIX	STANDARD	9	9	No

See Task 2 for parameter definitions.

Task 2

Display complete information on the client node named ARTHUR.

Command: Query Node arthur Format=Detailed

Output:

```

Node Name: ARTHUR
Platform: OS/2
Policy Domain Name: STANDARD
Last Access Date/Time: 05/19/1994 09:59:52
Days Since Last Access: 2
Password Set Date/Time: 05/18/1994 11:04:45
Days Since Password Set: 3
Locked?: No
Contact:
Compression: Client's Choice
Archive Delete Allowed?: Yes
Backup Delete Allowed?: No
Registration Date/Time: 05/18/1994 11:04:45
Registering Administrator: OPEN_REGISTRATION
Last Communication Method Used: Tcp/Ip
Bytes Received Last Session: 226
Bytes Sent Last Session: 556
Duration of Last Session (sec): 3.32
Pct. Idle Wait Last Session: 88.48
Pct. Comm. Wait Last Session: 6.63
Pct. Media Wait Last Session: 0.00
    
```

where:

Node Name

Specifies the name of the client node.

Platform

Specifies the platform of the client node. A question mark (?) is displayed until the client node first accesses the server and reports its platform type.

Policy Domain Name

Specifies the name of the policy domain to which the client node is assigned.

Last Access Date/Time

Specifies the last date and time that the client node accessed the server.

Days Since Last Access

Specifies the number of days that have elapsed since the last time the client node accessed the server.

Password Set Date/Time

Specifies the date and time that the client node's password was set.

Days Since Password Set

Specifies the number of days that have elapsed since the client node's password was set.

Locked?

Specifies whether the client node is locked out of ADSM.

Contact

Specifies any contact information for the client node.

Compression

Specifies whether compression is enabled.

Archive Delete Allowed?

Specifies whether the client node can delete his or her own archive files.

Backup Delete Allowed?

Specifies whether the client node can delete his or her own backup files.

Registration Date/Time

Specifies the date and time that the client node was registered.

Registering Administrator

Specifies the name of the administrator that registered the client node.

Last Communication Method Used

Specifies the communication method that was last used by the client node.

Bytes Received Last Session

Specifies the number of bytes received by the server during the last client node session.

Bytes Sent Last Session

Specifies the number of bytes sent during the last client node session.

Duration of Last Session (sec)

Specifies how long the most recent client node session lasted.

Pct. Idle Wait Last Session

Specifies, as a percentage of the total session time, the amount of time that the client was not performing any functions.

Pct. Comm. Wait Last Session

Specifies, as a percentage of the total session time, the amount of time that the client waited for a communication response from the server.

Pct. Media Wait Last Session

Specifies, as a percentage of the total session time, the amount of time that the client waited for a removable volume to be mounted.

Related Commands

Table 98 provides references to related commands.

Table 98. Commands Related to QUERY NODE

Command	Description	Reference
LOCK NODE	Prevents a client from accessing the server.	Page 245
REGISTER NODE	Defines a client to the server and sets options for that user.	Page 374
REMOVE NODE	Removes a client from the list of registered nodes for a specific policy domain.	Page 378
RENAME NODE	Changes the name for an existing user that belongs to a specific policy domain.	Page 384
UNLOCK NODE	Enables a locked user in a specific policy domain to access the server.	Page 436
UPDATE NODE	Changes the attributes associated with a client node.	Page 464

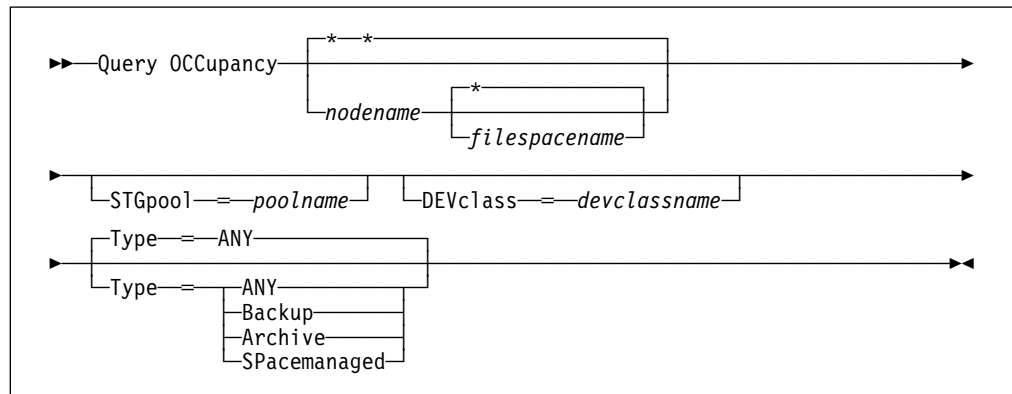
QUERY OCCUPANCY (Display File Space Information by Storage Pool)

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.

Privilege Class

Any administrator can issue this command.

Syntax



Parameters

nodename

Specifies the name of the node that owns the file spaces that you want to locate. This parameter is optional. You can use a pattern-matching expression to specify this name. All matching names are included in the query. If you do not specify a value for this parameter, all nodes are queried.

filespace name

Specifies the name of the file space that you want to locate. This parameter is optional. You can use a pattern-matching expression to specify this name. All matching names are included in the query. If you do not specify a value for this parameter, all file spaces are queried. If you explicitly specify a file space name, you must also specify a node name.

STGpool=poolname

Specifies the name of the storage pool to which the file spaces are assigned. This parameter is optional. You can use a pattern-matching expression to specify this name. All matching storage pools are included in the query. If you do not specify a value for this parameter, all storage pools are queried.

DEVclass=devclassname

Specifies the name of the device class to which the file spaces are assigned. This parameter is optional. You can use a pattern-matching expression to specify this name. All matching device classes are included in the query. If you do not specify a value for this parameter, all storage pools are queried.

Type=filetype

Specifies the storage type of the file spaces to be queried. This parameter is optional. The default value is ANY. Possible values are:

ANY

Specifies that all types of files are queried: backup versions of files, archived copies of files, and files migrated from one or more client nodes.

Backup

Specifies that backup file spaces are queried.

Archive

Specifies that archive file spaces are queried.

SPacemanaged

Specifies that space management files (files migrated from one or more client nodes) are queried.

Examples

Task 1

Display information on where all file spaces assigned to the node named BILLK are stored.

Command: Query OCCupancy billk

Output:

Node Name	Filespace Name	Storage Pool Name	Number of Files	Space Occupied (MB)
BILLK	IBMDOS_5	POOL2	21	1.02
BILLK	WP_FILES	POOLA	33	7.60
BILLK	SP_DATA	BACKUPPOOL	2	0.80

See Task 2 for parameter definitions.

Task 2

Display information on where the file spaces that belong to the node WAYNE are stored.

QUERY OCCUPANCY

Command: Query OCCupancy wayne

Output:

Node Name	Filespace Name	Storage Pool Name	Number of Files	Space Occupied (MB)
WAYNE	DWG1	BACKUPPOOL1	2,330	53.19
WAYNE	OS2C	BACKUPPOOL1	1,554	32.00
WAYNE	OS2D	BACKUPPOOL1	1,716	35.30

where:

Node Name

Specifies the name of the node that owns the specified file space.

Filespace Name

Specifies the name of the file space that belongs to the node.

Storage Pool Name

Specifies the name of the storage pool where the file space currently resides.

Number of Files

Specifies the number of files that belong to the file space and are stored in this storage pool.

Space Occupied (MB)

Specifies the amount of space occupied by the file space.

Related Commands

Table 99 provides references to related commands.

Table 99. Commands Related to QUERY OCCUPANCY

Command	Description	Reference
DELETE FILESPACE	Deletes data associated with one or more client's filespaces.	Page 167
QUERY FILESPACE	Displays information about data in one or more filespaces that belong to a client.	Page 300
QUERY NODE	Displays partial or complete information about one or more clients.	Page 320

QUERY OPTION (Query One or More Server Options)

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


The options file is the only place in which the server options can be changed. To change the options, you must restart the server.

For more information about the server options, see the *ADSM Installing the Server and Administrative Client*.

Privilege Class

Any administrator can issue this command.

Syntax



Parameters

optionname

Specifies the name of an option in the server options file for which you want information. This parameter is optional. You can use a pattern-matching expression to specify this name. All matching server options are displayed. If you do not specify this parameter, information on all options is displayed.

Examples

Task 1

Display general information about all defined server options.

Command: Query OPTion

Output:

Server Option	Option Setting	Server Option	Option Setting
CommTimeOut	600	IdleTimeOut	30
BufPoolSize	1024	LogPoolSize	128
DateFormat	1 (mm/dd/yyyy)	TimeFormat	1 (hh:mm:ss)
NumberFormat	1 (1,000.00)	MessageFormat	1
Language	AMENG	MaxSessions	25
ExpInterval	0	MirrorRead DB	Verify
MirrorRead LOG	Verify	MirrorWrite DB	Sequential
MirrorWrite LOG	Sequential	VolumeHistory	'adsmm2k.vol- ume.history'
Devconfig	ADSM2K.DEVFILE	TxnGroupMax	16
CommOpenTimeOut	20	TcpPort	1801
TcpName	TCPIP	IcsPort	1500
IcsSname	*NONE*	MsgSuppress	*NONE*
LuName	*NONE*	Iucv	*NONE*
RouteCode	11	MsgHighlight	*NONE*

QUERY OPTION

where:

Server Option

Specifies the name of the server option.

Option Setting

Specifies the current value defined to the server option.

Task 2

View the option settings for all options that begin with L.

Command: Query OPTion L*

Output:

Server Option	Option Setting
LogPoolSize	128
Language	AMENG
LuName	*NONE*

where:

Server Option

Specifies the name of the server option.

Option Setting

Specifies the current value defined to the server option.

Related Commands

None.

QUERY POLICYSET (Query a Policy Set)

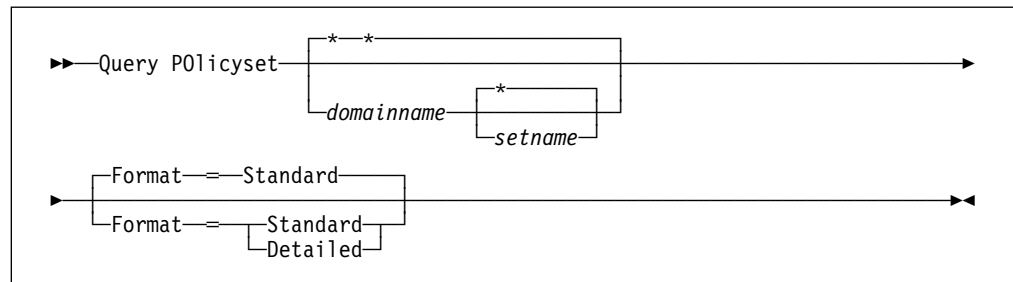
Use the QUERY POLICYSET command to display information about an existing policy set.

To get a list of the policy sets for a policy domain, use the QUERY DOMAIN command.

Privilege Class

Any administrator can issue this command.

Syntax



Parameters

domainname

Specifies the name of the policy domain associated with the copy group to be queried. This parameter is optional. You can use a pattern-matching expression to specify this name. All matching policy domains are queried. If you do not specify a value for this parameter, all policy domains are queried. You must specify this parameter when querying an explicitly named policy set.

setname

Specifies the name of the policy set to be queried. This parameter is optional. You can use a pattern-matching expression to specify this name. All matching policy sets are queried. If you do not specify either ACTIVE or the policy set name, all policy sets are queried. You must specify this parameter when querying an explicitly named policy set.

Format=*formatvalue*

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed for the specified policy set.

Detailed

Specifies that complete information is displayed for the specified policy set.

Examples

Task 1

Query all policy sets for all policy domains. Create the output in standard format.

Command: Query Policyset

Output:

Policy Domain Name	Policy Set Name	Default Mgmt Class Name	Description
-----	-----	-----	-----
ENGDOMAIN	ACTIVE	STANDARD	Installed default policy set.
ENGDOMAIN	STANDARD	STANDARD	Installed default policy set.
MARIE	ACTIVE	DAVID	
MARIE	DAVIDTEST	DAVID	
STANDARD	ACTIVE	STANDARD	Installed default policy set.
STANDARD	STANDARD	STANDARD	Installed default policy set.
TANGDOMA- IN	ACTIVE	TANGMC3	This is my policy set!
TANGDOMA- IN	TANGPOLI- CY	TANGMC3	This is my policy set!

See Task 2 for parameter definitions.

Task 2

Query the VACATION policy set that is defined to the EMPLOYEE_RECORDS policy domain. Create the output in detailed format.

Command: Query Policyset employee_records vacation Format=Detailed

Output:

```

Policy Domain Name: EMPLOYEE_RECORDS
Policy Set Name: VACATION
Default Mgmt Class Name: ACTIVEFILES
Description: Programming Group
Last Update by (administrator): TOM
Last Update Date/Time: 04/26/1994 11:30:06
    
```

where:

Policy Domain Name

Specifies the name of the policy domain.

Policy Set Name

Specifies the name of the activated policy set for the policy domain.

Default Mgmt Class Name

Specifies the assigned default management class for the policy set.

Description

Specifies a description of the policy domain.

Last Update by (administrator)

Specifies the name of the administrator that defined or most recently updated the policy domain.

Last Update Date/Time

Specifies when the administrator defined or most recently updated the storage pool.

Related Commands

Table 100 provides references to related commands.

Table 100. Commands Related to QUERY POLICYSET

Command	Description	Reference
ACTIVATE POLICYSET	Validates and activates a policy set within a policy domain.	Page 38
COPY POLICYSET	Creates a copy of an existing policy set.	Page 75
DEFINE POLICYSET	Defines a policy set within the specified policy domain.	Page 118
DELETE POLICYSET	Deletes a policy set, including its management classes and copy groups, from a policy domain.	Page 174
QUERY DOMAIN	Displays information about one or more policy domains.	Page 287
UPDATE POLICYSET	Changes the description of a policy set.	Page 467
VALIDATE POLICYSET	Verifies and reports on conditions the administrator must consider before activating the policy set.	Page 501

QUERY PROCESS (Query One or More Server Processes)

Use the QUERY PROCESS command to display information about one or more active background processes. The following commands can generate background processes:

AUDIT LICENSES	EXPORT SERVER
AUDIT VOLUME	EXTEND LOG
BACKUP DB	IMPORT ADMIN
BACKUP STGPOOL	IMPORT NODE
DEFINE DBCOPY	IMPORT POLICY
DEFINE LOGCOPY	IMPORT SERVER
DELETE DBVOLUME	MOVE DATA
DELETE FILESPACE	REDUCE DB
DELETE LOGVOLUME	REDUCE LOG
DELETE VOLUME	RESTORE STGPOOL
EXPIRE INVENTORY	RESTORE VOLUME
EXPORT ADMIN	VARY
EXPORT NODE	
EXPORT POLICY	

The following internal server operations generate background server processes:


- Inventory expiration
- Migration
- Reclamation

You can cancel background processes, by process number, using the CANCEL PROCESS command.

Privilege Class

Any administrator can issue this command.

Syntax



Parameters

processnum

Specifies the number of the background process to be queried. This parameter is optional. If you do not specify a process number, information on all background processes is displayed.

Examples

Task 1

Display information about background process 2 that was initiated with the MOVE DATA command.

Command: Query PRocess 2

Output:

Process Number	Process Description	Status
2	Move Data	Volume 0300 (storage pool BACKUPPOOL), Target Pool BACKUPPOOL, Moved Files: 20, Moved Bytes: 696320, Unreadable Files: 0, Unreadable Bytes: 0. Current File (bytes): None

where:

Process Number

Specifies the number assigned to the background process.

Process Description

Specifies a description of the background process.

Status

Specifies the completion status of the background process.

Task 2

Display information about all background processes.

Command: Query PProcess

Output:

Process Number	Process Description	Status
2	Move Data	Volume 0300 (storage pool BACKUPPOOL), Target Pool BACKUPPOOL, Moved Files: 20, Moved Bytes: 696320, Unreadable Files: 0, Unreadable Bytes: 0. Current File (bytes): None
3	Extend Log	9 megabytes of 12 formatted.

where:

Process Number

Specifies the number assigned to the background process.

Process Description

Specifies a description of the background process.

Status

Specifies the completion status of the background process.

QUERY PROCESS

Related Commands

Table 101 provides a reference to a related command.

Table 101. Command Related to QUERY PROCESS

Command	Description	Reference
CANCEL PROCESS	Cancels a background server process.	Page 64

QUERY SCHEDULE (Query One or More Schedules)

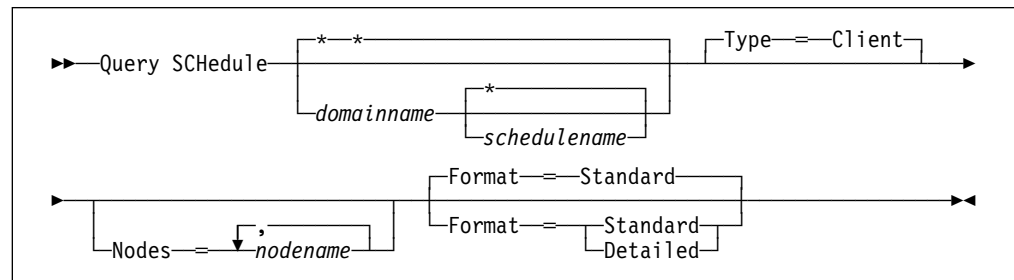
Use the QUERY SCHEDULE command to display information pertaining to one or more schedules.

The QUERY SCHEDULE command takes two forms, depending on whether the schedule applies to client operations or administrative commands. The syntax and parameters for each form are defined separately.

Privilege Class

Any administrator can issue this command.

Syntax for Querying Client Schedules



Parameters for Querying Client Schedules

domainname

Specifies the name of the policy domain to which the schedule belongs. You can use a pattern-matching expression to specify this name. All matching policy domains are queried. The default queries all existing policy domains. If you specify a domain name, you do not have to specify a schedule name.

schedulename

Specifies the name of the schedule that belongs to the specified policy domain. You can use a pattern-matching expression to specify this name. All matching schedules are queried. The default queries all existing client schedules for a given domain. If you specify a schedule name, you must also specify a policy domain name.

Type=Client

Specifies that this command queries client schedules. This parameter is optional. The default value is CLIENT. Administrative command schedules are queried by specifying TYPE=ADMINISTRATIVE.

Nodes=nodename

Specifies that you want to restrict the schedules that are queried by defining one or more client nodes that are associated with the schedules to be displayed. This parameter is optional. You can use a pattern-matching expression to specify a name. All matching clients are queried. If you do not specify a client name, all schedules matching the DOMAINNAME and SCHEDULENAME parameters are displayed. You can specify more than one client node name by separating each name with a comma, with no intervening spaces.

QUERY SCHEDULE

Format=*formatvalue*

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:


Standard

Specifies that partial information is displayed for the specified schedules.

Detailed

Specifies that complete information is displayed for the specified schedules.

Syntax for Querying Administrative Command Schedules



Parameters for Querying Administrative Command Schedules

schedulename

Specifies the name of the schedule to be queried. You can use a pattern-matching expression to specify this name. All matching schedules are queried. The default queries all existing administrative schedules.

Type=Administrative

Specifies that the schedules to be queried are administrative schedules. This parameter is required.

Format=*formatvalue*

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed for the specified schedules.

Detailed

Specifies that complete information is displayed for the specified schedules.

Examples

Task 1

Display all schedules that belong to the PROG1 policy domain.

Command: Query SCHEDULE prog1

Output:

Domain	* Schedule Name	Action	Start Date/Time	Duration	Period	Day
MON	TOM	Inc Bk	03/09/1995 12:55:18	1 H	1 D	Any

where:

Domain

Specifies the name of the policy domain to which the specified schedule belongs.

*** (Asterisk)**

Specifies whether the corresponding schedule has expired. If there is an asterisk in this column, the corresponding schedule has expired.

Schedule Name

Specifies the name of the schedule.

Action

Specifies the action that occurs when this schedule executes.

Start Date/Time

Specifies the initial starting date and time for this schedule.

Duration

Specifies the length of the startup window.

Period

Specifies the time between startup windows (assuming DAYOFWEEK=ANY).

Day

Specifies the day of the week on which the startup windows begin.

Task 2

Display, in detail, the administrative command schedule (BACKUP_ARCHIVEPOOL) defined to backup the primary storage pool ARCHIVEPOOL to the copy storage pool RECOVERYPOOL.

Command: Query SCHEDULE backup_archivepool Type=Administrative
Format=Detailed

Output:

```

Schedule Name: BACKUP_ARCHIVEPOOL
Description: backup data in primary storage pool
Options:
Command: Backup STGpool archivepool recoverypool
Priority: 5
Start Date/Time: 03/28/1995 16:20:48
Duration: 1 Hour(s)
Period: 2 day(s)
Day of Week: Monday
Expiration:
Last Update by (administrator): JANET
Last Update Date/Time: 03/28/1995 16:35:06
Active?: No
    
```

where:

Schedule Name

Specifies the name of the schedule.

QUERY SCHEDULE

Description

Specifies the description of the schedule.

Options

Specifies the options supplied to the DSMC command at the time the schedule is executed.

Command

Specifies the command that is scheduled.

Priority

Specifies the priority value for this schedule.

Start Date/Time

Specifies the initial starting date and time for this schedule.

Duration

Specifies the length of the startup window.

Period

Specifies the time between startup windows (assuming DAYOFWEEK=ANY).

Day of Week

Specifies the day of the week on which the startup windows begin.

Expiration

Specifies the date and time on which this schedule will no longer be used. If this column is blank, the schedule does not expire.

Last Update by (administrator)

Specifies the name of the administrator that most recently modified the schedule.

Last Update Date/Time

Specifies the last date and time the schedule was modified.

Active?

Specifies whether the schedule has been processed according to the time and date set for this schedule.

Related Commands

Table 102 provides references to related commands.

Table 102. Command Related to QUERY SCHEDULE

Command	Description	Reference
COPY SCHEDULE	Creates a copy of an existing schedule.	Page 77
DEFINE SCHEDULE	Defines a schedule for a client operation or an administrative command.	Page 121
UPDATE SCHEDULE	Changes the attributes of a schedule.	Page 469


QUERY SESSION (Query One or More Client Sessions)

Use the QUERY SESSION command to display information on one or more administrative and client node sessions.

Privilege Class

Any administrator can issue this command.

Syntax



Parameters

sessnum

Specifies the number of the administrative or client node session to be queried. This parameter is optional. If you do not specify a value for this parameter, all sessions are displayed.

Format=*formatvalue*

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed for the session.

Detailed

Specifies that complete information is displayed for the session.

Examples

Task 1

Display information on all administrative and client node sessions that are communicating with the server.

Command: Query SEssion

Output:

Sess Number	Comm. Method	Sess State	Wait Time	Bytes Sent	Bytes Recvd	Sess Type	Platform	Client Name
2	Tcp/Ip	Run	0 S	5.4 K	735	Admin	OS/2	HAYDEN

See Task 2 for parameter definitions.

Task 2

Display information in full detail about all administrative and client node sessions that are communicating with the server.

Command:Query SSession Format=Detailed

Output:

```
Sess Number: 20
Comm. Method: Tcp/Ip
Sess State: MediaW
Wait Time: 36.2 M
Bytes Sent: 165.2 K
Bytes Recvd: 112.9 K
Sess Type: Node
Platform: OS/2
Client Name: JSBACH
Media Access Status: Waiting for mount point in device
                    class TAPDEV5 (2170 seconds).
```

where:

Sess Number

Specifies a unique session identification number assigned by the server.

Comm. Method

Specifies the method being used by the client to communicate with the server.

Sess State

Specifies the current communications state of the server. Possible states are:

- End** The session is ending (session resources being released).
- IdleW** Waiting for client's next request (session is *idle*).
- MediaW** The session is waiting for access to a sequential access volume.
- RecvW** Waiting to receive an expected message from the client.
- Run** The server is executing a client request (and not waiting to send data).
- SendW** The server is waiting to send data to the client (waiting for data to be delivered to the client node that has already been sent).
- Start** The session is starting (authentication is in progress).

Wait Time

Specifies the amount of time (Seconds, Minutes, or Hours) the server has been in the current state shown.

Bytes Sent

Specifies the number of bytes of data sent to the client node since the session was initiated.

Bytes Recvd

Specifies the number of bytes of data received from the client node since the session was initiated.

Sess Type

Specifies the type of session in process: either ADMIN for an administrative session, or NODE for a client node session.

Platform

Specifies the type of operating system associated with the client.

Client Name

Specifies the name of the administrator or client node.

Media Access Status

Specifies the type of media wait state. This information is only provided when the session state equals *MediaW*.

Related Commands

Table 103 provides a reference to a related command.

Table 103. Command Related to QUERY SESSION

Command	Description	Reference
CANCEL SESSION	Cancels one or all active sessions with the server.	Page 66

QUERY STATUS (Query System Parameters)

Use the QUERY STATUS command to display information about the general server parameters, such as those defined by the SET commands.

You can use the QUERY STATUS command to request information about client sessions, such as the availability of the server, password authentication, accounting settings, or the retention period for the information retained in the activity log. You can also use this command to display information about the central scheduler, such as the central scheduling mode of the server, or to display the maximum number of retries allowed after a failed attempt to execute a scheduled command. Also available is information about the license audit period, and the last date and time that a license audit was performed on the server.

Privilege Class

Any administrator can issue this command.

Syntax

```
▶▶—Query SStatus—————▶▶
```

Parameters

None.

Examples

Task

Display general information about server parameters.

Command: Query SStatus

Output:

```

Server Name: ADSM
Server Installation Date/Time: 04/12/1995 14:06:20
Server Restart Date/Time: 04/21/1995 09:23:00
Authentication: Off
Password Expiration Period: 90 Day(s)
Registration: Closed
Availability: Enabled
Accounting: Off
Activity Log Retention Period: 1 Day(s)
License Audit Period: 30 Day(s)
Last License Audit: 04/12/1995 14:07:43
Server License Compliance: Valid
Central Scheduler: Active
Maximum Sessions: 25
Maximum Scheduled Sessions: 12
Event Record Retention Period: 10 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: Enabled
    
```

where:

Server Name

Specifies the name of the server.

Server Installation Date/Time

Specifies the date and time the server was installed.

Server Restart Date/Time

Specifies the last date and time the server was started.

Authentication

Specifies whether password authentication is set on or off.

Password Expiration Period

Specifies the period, in days, after which the administrator or client node password expires.

Registration

Specifies whether client node registration is open or closed.

Availability

Specifies whether the server is enabled or disabled.

Accounting

Specifies whether an accounting record is generated at the end of each client node session.

Activity Log Retention Period

Specifies the number of days information is retained in the activity log.

License Audit Period

Specifies the period in days between automatic license audits that are performed by the ADSM server. Additional licensing information is available by using the QUERY LICENSE command.

Last License Audit

Specifies the date and time the last license audit was performed on the ADSM server. Additional licensing information is available by using the QUERY LICENSE command.

Server License Compliance

Indicates whether the server is in compliance with license terms. Possible values are:

- Valid—server complies with license terms
- Failed—server does *not* comply with license terms

To see what factors are causing the server to fail to comply with license terms, use the QUERY LICENSE command.

Central Scheduler

Specifies whether central scheduling is running (active or aborted).

Maximum Sessions

Specifies the maximum number of client/server sessions.

Maximum Scheduled Sessions

Specifies the maximum number of client/server sessions available for processing scheduled work.

Event Record Retention Period

Specifies the number of days central scheduler event records are retained.

Schedule Randomization Percentage

Specifies how much of the startup window is used for executing scheduled events in client-polling mode.

Query Schedule Period

Specifies the frequency with which clients poll the server to obtain scheduled work, in client-polling mode. If the value in this field is *Client's Choice*, the polling frequency is determined by the client node.

Maximum Command Retries

Specifies the maximum number of times that a client scheduler retries after a failed attempt to execute a scheduled command. If the value in this field is *Client's Choice*, the client node determines the maximum number of retries.

Retry Period

Specifies the number of minutes between attempts by the client scheduler to retry after a failed attempt to contact the server or to execute a scheduled command. If the value in this field is *Client's Choice*, the client node determines the number of minutes between retries.

Scheduling Modes

Specifies the central scheduling modes supported by the server.

Log Mode

Specifies the mode for saving recovery log records. The mode is either roll-forward or normal.

Database Backup Trigger

Specifies that ADSM automatically runs a backup of the database if the backup trigger is enabled.

Related Commands

Table 104 provides references to related commands.

Table 104 (Page 1 of 2). Commands Related to QUERY STATUS

Command	Description	Reference
DISABLE	Prevents new clients from accessing ADSM but permits existing sessions to continue.	Page 188
ENABLE	Resumes server activity following the DISABLE command.	Page 192
QUERY LICENSE	Displays information about licenses and audits.	Page 304
SET ACCOUNTING	Specifies whether accounting records are created at the end of each client session.	Page 407
SET ACTLOGRETENTION	Specifies the number of days to retain log records in the activity log.	Page 409
SET AUTHENTICATION	Specifies whether clients are prompted for a password at the beginning of a session.	Page 411
SET EVENTRETENTION	Specifies the number of days to retain event records.	Page 412
SET LOGMODE	Sets the mode for saving records in the recovery log.	Page 416
SET MAXCMDRETRIES	Specifies the maximum number of retries after a failed attempt to execute a scheduled command.	Page 419
SET MAXSCHEDSESSIONS	Specifies the maximum number of client/server sessions available for processing scheduled work.	Page 421
SET PASSEXP	Specifies the number of days after which a password is expired and must be changed.	Page 423

QUERY STATUS

Table 104 (Page 2 of 2). Commands Related to QUERY STATUS

Command	Description	Reference
SET QUERYSCHEDPERIOD	Specifies the frequency for clients to obtain scheduled work, in client-polling mode.	Page 424
SET RANDOMIZE	Specifies the randomization of start times within a window for schedules in client-polling mode.	Page 426
SET REGISTRATION	Specifies whether users can register themselves or must be registered by an administrator.	Page 428
SET RETRYPERIOD	Specifies the time between retry attempts by the client scheduler.	Page 430
SET SCHEDMODES	Specifies the central scheduling mode for the server.	Page 432
SET SERVERNAME	Specifies the name by which the server is identified.	Page 434

QUERY STGPOOL (Query One or More Storage Pools)

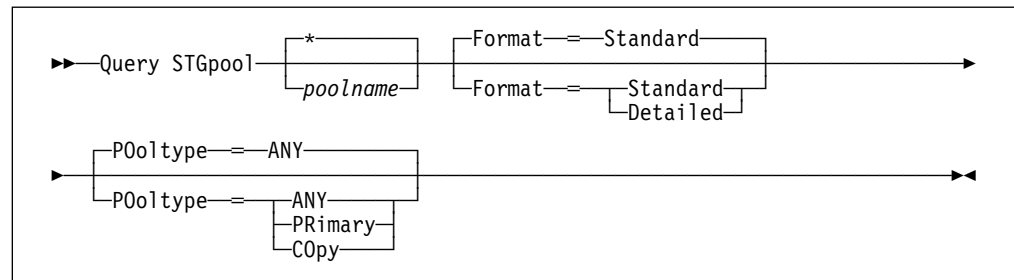
Use the QUERY STGPOOL command to display information about one or more storage pools.

You can also use this command to monitor migration processes for the storage pool.

Privilege Class

Any administrator can issue this command.

Syntax



Parameters

poolname

Specifies the name of the storage pool to be queried. This parameter is optional. You can use a pattern-matching expression to specify this name. All matching storage pools are included. If you do not specify a value for this parameter, all storage pools are displayed.

Format=*formatvalue*

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed for the storage pool.

Detailed

Specifies that complete information is displayed for the storage pool.

Pooltype=*type*

Specifies the type of storage pool to be queried. This parameter is optional. The default value is ANY. Possible values are:

ANY

Specifies that primary and copy storage pools will be queried.

PRimary

Specifies that only primary storage pools will be queried.

COpy

Specifies that only copy storage pools will be queried.

Usage Notes

- For storage pools defined to DISK devices, if CACHE=YES, the value for %UTIL is consistently high, even after migration, as the device's capacity is maximized with cached copies of data. The value for %MIGR drops after migration, indicating that less data is available to be migrated.
- For storage pools defined to DISK devices, the value for %UTIL can be higher than the value for %MIGR if you issue this command while a file creation transaction is in progress. The value for %UTIL is determined by the amount of space actually allocated (while the transaction is in progress) and the value for %MIGR only represents the space occupied by *committed* files. At the end of the transaction, these values become synchronized.

Examples

Task 1

Display information for the storage pool named POOL1.

Command: Query STGpool pool1

Output:

Storage Pool Name	Device Class Name	Estimated Capacity (MB)	%Util	%Migr	High Mig%	Low Mig%	Next Storage Pool
POOL1	DISK	58.5	0.8	0.7	90	70	POOL2

See Task 2 for parameter definitions.

Task 2

Display information in full detail for the storage pool named POOL1.

Command: Query STGpool pool1 Format=Detailed

Output:

```

Storage Pool Name: POOL1
Storage Pool Type:
Device Class Name: DISK
Estimated Capacity (MB): 58.5
                    %Util: 0.0
                    %Migr: 0.0
                    High Mig%: 90
                    Low Mig%: 70
Migration Processes:
Next Storage Pool: SPOOL2
Maximum Size Threshold: No Limit
Access: Read/Write
Description:
Cache Migrated Files?: Yes
Collocate?:
Reclamation Threshold:
Maximum Scratch Volumes Allowed:
Delay Period for Volume Reuse:
Migration in Progress?:
Amount Migrated (MB):
Elapsed Migration Time (seconds):
Reclamation in Progress?:
Volume Being Migrated/Reclaimed:
Last Update by (administrator): HTANG
Last Update Date/Time: 04/07/1995 12:37:05

```

where:

Storage Pool Name

Specifies the name of the storage pool.

Storage Pool Type

Specifies the type of storage pool you want to query.

Device Class Name

Specifies the name of the device class assigned to the storage pool.

Estimated Capacity (MB)

Specifies the estimated capacity of the storage pool, in megabytes.

For DISK devices, this is the capacity of all volumes in the storage pool, including any volumes that are varied offline.

For sequential access devices, this is an estimate of the total space of all the sequential access volumes in the storage pool, including volumes that are *unavailable* or *readonly* and all volumes defined with the MAXSCRATCH parameter.

%Util

Specifies, as a percentage, an estimate of the utilization of the storage pool.

For DISK devices, this is specified as a percentage of the value for the estimated capacity, including cached data and data that resides on any volumes which are varied offline.

For sequential access devices, this is specified as a percentage of the number of active bytes on each sequential access volume and the estimated capacity of all volumes in the storage pool.

The %UTIL value includes cached data on a volume (when cache is enabled); the %MIGR value excludes cached data. Therefore, when cache is enabled and migration occurs, the %MIGR value decreases but the %UTIL value remains the same because the migrated data is remaining on the volume as cached data. In this instance, the %UTIL value decreases only when the cached data expires or when the space that cached files are using needs to be used for noncached files.

%Migr

Specifies an estimate of the percentage of data in the storage pool that can be migrated. This value is used by the high and low migration thresholds to determine when to start and stop migration.

For DISK devices, this value is specified as a percentage of the value for the estimated capacity, excluding cached data, but including data on any volumes varied offline.

For sequential access devices, this value is specified as the percentage of the number of sequential access volumes in the storage pool that contain at least 1 byte of active data.

The %UTIL value includes cached data on a volume (when cache is enabled); the %MIGR value excludes cached data. Therefore, when cache is enabled and migration occurs, the %MIGR value decreases but the %UTIL value remains the same because the migrated data is remaining on the volume as cached data. In this instance, the %UTIL value decreases only when the cached data expires or when the space that cached files are using needs to be used for noncached files.

High Mig%

Specifies when ADSM can begin migration for the storage pool. Migration processes can begin when the percentage of data that can be migrated reaches this threshold.

Low Mig%

Specifies when ADSM can stop migration for the storage pool. Migration processes can stop when the percentage of data on the volume that can be migrated reaches this threshold.

Migration Process

For random access storage pools only, specifies the number of processes that are used for migrating files from a storage pool. Each migration process uses a drive so the number of migration processes should not exceed the number of drives that can be dedicated to migration processing.

Next Storage Pool

Specifies the destination storage pool for data that is migrated from this storage pool.

Maximum Size Threshold

Specifies the maximum file size for files in the storage pool.

Access

Specifies how the data in the storage pool can be accessed.

Description

Specifies a description of the storage pool.

Cache Migrated Files?

For DISK device classes only, specifies whether caching is enabled.

Collocate?

For sequential access device classes, specifies whether collocation is enabled.

Reclamation Threshold

Specifies when data from volumes in a storage pool is reclaimed. Each individual volume's percentage of reclaimable space is compared to this value to determine if reclamation is necessary.

Maximum Scratch Volumes Allowed

Specifies the maximum number of scratch volumes that can be requested by ADSM.

Delay Period for Volume Reuse

Specifies the number of days that must elapse after files have been deleted from a volume, before the volume can be returned to scratch status or reused.

Migration in Progress?

Specifies whether a migration process is active for the storage pool.

Amount Migrated (MB)

Specifies the amount of data, in megabytes, that has been migrated, if a migration process is active. If no migration is in progress, this value indicates the amount of data migrated during the last migration process.

Elapsed Migration Time (seconds)

Specifies the amount of time that has passed since the process started, if a migration process is active. If no migration is in progress, this value indicates the amount of time required to complete the last migration process.

Reclamation in Progress?

Specifies whether a reclamation process is active for the storage pool.

Volume Being Migrated/Reclaimed

For sequential access storage pools, specifies the name of a sequential access volume for which migration or reclamation processing is active. If no migration or reclamation processes are active on a volume, this field is blank. For random access storage pools, this field is always blank.

Last Update by (administrator)

Specifies the name of the administrator that has defined or most recently updated the storage pool.

Last Update Date/Time

Specifies when the administrator has defined or most recently updated the storage pool.

Related Commands

Table 105 provides references to related commands.

Table 105. Commands Related to QUERY STGPOOL

Command	Description	Reference
DEFINE STGPOOL	Defines a storage pool as a named collection of data storage media.	Page 135
DELETE STGPOOL	Deletes a storage pool from data storage.	Page 179
UPDATE STGPOOL	Changes the attributes of a storage pool.	Page 482

QUERY VOLHISTORY (Display Sequential Volume History Information)

Use the QUERY VOLHISTORY command to display sequential volume history information that has been collected by the server.

Volume history information is collected about the following types of sequential volumes accessed by the server:

- Volumes used to back up an ADSM database (full or incremental)
- Volumes used to dump an ADSM database
- Volumes used to export administrator, node, policy, or server data
- Storage pool sequential access volumes that have been added
- Storage pool sequential access volumes that have been reused in storage pools through reclamation or MOVE DATA operations
- Storage pool sequential access volumes that have been removed by using the DELETE VOLUME command or removed during reclamation of scratch volumes

Sequential access volumes are removable media volumes that are accessed sequentially by the ADSM server, such as removable tape or optical volumes.

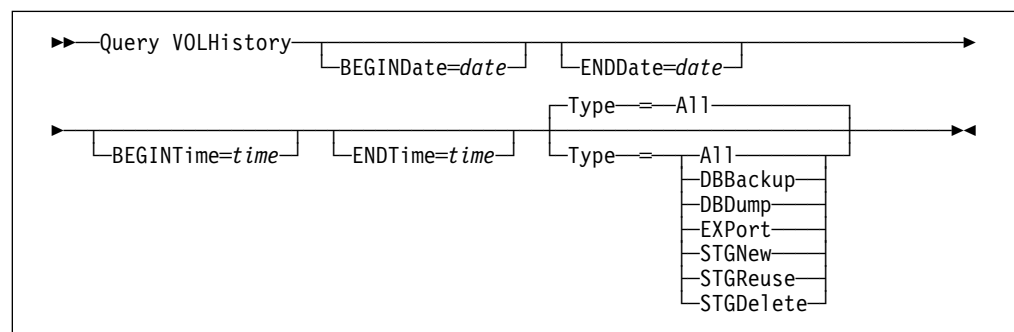
To save sequential volume history information collected by the server to one or more files, see the BACKUP VOLHISTORY command.

To delete sequential volume history information collected by the server, use the DELETE VOLHISTORY command.

Privilege Class

Any administrator can issue this command.

Syntax



Parameters

BEGINDate=date

Specifies the beginning date for the history information to display. This parameter is optional. The default is the earliest date for which history information has been collected. Possible values are:

MM/DD/YYYY

Specifies that you want to display information beginning with records created on a specific date, using MM/DD/YYYY format (for example, 01/31/1995).

TODAY

Specifies that you want to display information beginning with records created on the current date.

TODAY-*numdays* or -*numdays*

Specifies that you want to display information beginning with records created on the current date minus a specific number of days. The maximum value for the number of days is 9999.

For example, to display information beginning with records created a week ago, you can specify TODAY-7 or simply -7.

ENDDate=*date*

Specifies the ending date for the history information to display. This parameter is optional. The default is the current date. Possible values are:

MM/DD/YYYY

Specifies that you want to display information ending with records created on a specific date using MM/DD/YYYY format (for example, 02/28/95).

TODAY

Specifies that you want to display information ending with records created on the current date.

TODAY-*numdays* or -*numdays*

Specifies that you want to display information ending with records created on or before the current date minus a specific number of days. The maximum value for the number of days is 9999.

For example, to display all volume history information created up to yesterday, you can specify TODAY-1 or simply -1.

BEGINTime=*time*

Specifies the beginning time for the history information to display. This parameter is optional. The default is midnight (00:00:00) on the date specified with the BEGINDATE parameter. Possible values are:

HH:MM:SS

Specifies that you want to display information beginning with records created at a specific time on the specified begin date using HH:MM:SS format (for example, 12:33:28).

NOW

Specifies that you want to display information beginning with records created at the current time on the specified begin date.

NOW+*numhours:numminutes* or +*numhours:numminutes*

Specifies that you want to display information beginning with records created at the current time plus a specific number of hours (and, optionally, minutes) on the specified begin date.

For example, if you issue the QUERY VOLHISTORY command at 9:00 with `BEGINTIME=NOW+3` or `BEGINTIME=+3`, ADSM displays records with a time of 12:00 or later on the specified begin date.

NOW–numhours:numminutes or –numhours:numminutes

Specifies that you want to display information beginning with records created at the current time minus a specific number of hours (and optionally minutes) on the specified begin date.

For example, if you issue the QUERY VOLHISTORY command at 9:00 with `BEGINTIME=NOW–3:30` or `BEGINTIME=–3:30`, ADSM displays records with a time of 5:30 or later on the specified begin date.

ENDTime=time

Specifies the ending time for the history information to display. This parameter is optional. The default is the current time. Possible values are:

HH:MM:SS

Specifies that you want to display information ending with records created at a specific time on the specified end date using HH:MM:SS format (for example, 18:59:59).

NOW

Specifies that you want to display information ending with records created at the current time on the specified end date.

NOW+numhours:numminutes or +numhours:numminutes

Specifies that you want to display information ending with records created a specific number of hours (and optionally, minutes) after the current time on the specified end date.

For example, if you issue the QUERY VOLHISTORY command at 9:00 with `ENDTIME=NOW+3` or `ENDTIME=+3`, ADSM displays records with a time of 12:00 or earlier on the specified end date.

NOW–numhours:numminutes or –numhours:numminutes

Specifies that you want to display information ending with records created a specific number of hours (and, optionally, minutes) prior to the current time on the specified end date.

For example, if you issue the QUERY VOLHISTORY command at 9:00 with `ENDTIME=NOW–3:30` or `ENDTIME=–3:30`, ADSM displays records with a time of 5:30 or earlier on the specified end date.

Type=typevalue

Specifies the type of records to display from the volume history file. This parameter is optional. The default is ALL. Possible values are:

All

Specifies that all history information is displayed.

DBBackup

Specifies to display only records that contain information about volumes used for backups (full or incremental).

DBDump

Specifies to display only records that contain information about volumes used for online database dumps (done before installation of Version 2 Release 1).

QUERY VOLHISTORY

EXPort

Specifies to display only records that contain information about volumes used for exporting data.

STGNew

Specifies to display only records that contain information about new sequential access storage volumes.

STGReuse

Specifies to display only records that contain information about reused sequential storage pool volumes.

STGDelete

Specifies to display only records that contain information about deleted sequential storage pool volumes.

Examples

Task

Display all volume history information created a week before March 31, 1995.

Command: Query VOLHistory BeGINdate=-7

Output:

```
      Date/Time: 03/24/1995 05:30:06
      Volume Type: STGNEW
      Backup Series: 0
      Backup Operation: 0
      Volume Seq: 0
      Device Class: DC480
      Volume Name: 8MNL65
      Volume Location:

      Date/Time: 03/24/1995 05:30:44
      Volume Type: STGDELETE
      Backup Series: 0
      Backup Operation: 0
      Volume Seq: 0
      Device Class: DC480
      Volume Name: 8MNL65
      Volume Location:
```

where:

Date/Time

The date and time for the volume information.

Volume Type

The type of volume. Possible values are:

BACKUPFULL

The volume was used for a full backup of the database.

BACKUPINCR

The volume was used for an incremental backup of the database.

DBDUMP

The volume was used for an online dump of the database. (These volumes were used before the installation of ADSM Version 2 Release 1.)

EXPORT

The volume was used to record data for export.

STGDELETE

The volume was a sequential access storage pool volume that was deleted.

STGNEW

The volume was a sequential access storage pool volume that was added.

STGREUSE

The volume was a sequential access storage pool volume that was reused.

Backup Series

If the volume type is BACKUPFULL or BACKUPINCR, the identifying number of the backup series of which this volume is a part. Each backup series consists of a full backup and all incremental backups that apply to that full backup, up to the next full backup of the database.

Backup Operation

If the volume type is BACKUPFULL or BACKUPINCR, the operation number of this backup volume, within the backup series. The full backup within a backup series is operation 0. The first incremental backup for that full backup is operation 1, the second incremental backup is operation 2, and so on.

Volume Seq

If the volume type is BACKUPFULL, BACKUPINCR, or EXPORT, the sequence number of the volume when it was used for backing up the database or for exporting data.

For BACKUPFULL and BACKUPINCR volumes, this is the sequence, or position, of the volume within the backup series. Volume sequence 1 identifies the first volume used for the first operation (a full backup), and so on. For example, if the full backup occupies three volumes, these volumes are identified as volume sequence 1, 2, and 3, respectively. The first volume of the next operation (the first incremental backup) is then volume sequence 4.

Device Class

The device class used for the volume.

Volume Name

The name of the volume.

Volume Location

The location of the volume. The information is available only for BACKUPFULL, BACKUPINCR, DBDUMP, and EXPORT volumes, and only if the UPDATE VOLHISTORY command was used to add the information.

Related Commands

Table 106 provides references to related commands.

Table 106. Commands Related to QUERY VOLHISTORY

Command	Description	Reference
BACKUP VOLHISTORY	Records sequential volume history information in external files.	Page 60
DELETE VOLHISTORY	Removes sequential volume history information from the volume history file.	Page 181
UPDATE VOLHISTORY	Adds or changes location information for a volume in the volume history file.	Page 493

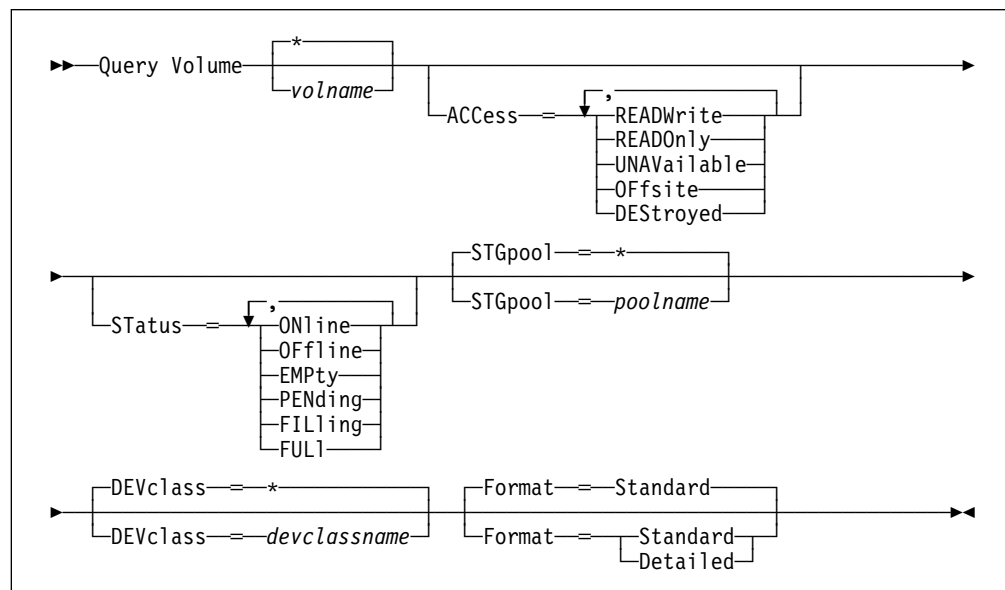
QUERY VOLUME (Query One or More Storage Pool Volumes)

Use the QUERY VOLUME command to display information about volumes in primary and copy storage pools. This command displays information about scratch volumes defined to ADSM.

Privilege Class

Any administrator can issue this command.

Syntax



Parameters

volname

Specifies the name of the volume to be displayed. This parameter is optional. You can use a pattern-matching expression to specify this name. All matching volumes are included in the query. If you do not specify a storage pool volume name, all storage pool volumes are included in the query.

ACCess=accessmode

Specifies that output should be restricted by volume access mode. This parameter is optional. You can specify more than one access mode in a list by delimiting each mode with a comma and no intervening spaces. If you do not specify a value for this parameter, output is not restricted by access mode. Possible values are:

READWrite

Specifies that volumes with an access mode of READWRITE are displayed.

READOnly

Specifies that volumes with an access mode of READONLY are displayed.

UNAVailable

Specifies that volumes with an access mode of UNAVAILABLE are displayed.

OFFsite

Specifies that volumes with an access mode of OFFSITE are displayed.

DESTroyed

Specifies that volumes with an access mode of DESTROYED are displayed.

Status=*status*

Specifies that output should be restricted by volume status. This parameter is optional. You can specify more than one status in a list by delimiting each status with a comma and no intervening spaces. If you do not specify a value for this parameter, output is not restricted by volume status. Possible values are:

ONline

Specifies that volumes with a status of ONLINE are displayed.

OFFline

Specifies that volumes with a status of OFFLINE are displayed.

EMPTy

Specifies that volumes with a status of EMPTY are displayed.

PENding

Specifies that volumes with a status of PENDING are displayed. These are volumes from which all files have been deleted, but the time specified by the REUSEDELAY parameter from the DEFINE STGPOOL command has not elapsed.

FILLing

Specifies that volumes with a status of FILLING are displayed.

FULI

Specifies that volumes with a status of FULL are displayed.

STGPool=*poolname*

Specifies the name of the storage pool to be included in the query. This parameter is optional. You can use a pattern-matching expression to specify this name. All matching storage pools are included in the query. If you do not specify a storage pool name, all storage pools are included in the query.

DEVclass=*devclassname*

Specifies the name of the device class to be included in the query. This parameter is optional. You can use a pattern-matching expression to specify this name. All matching device classes are included in the query. If you do not specify a device class name, all devices are included in the query.

Format=*formatvalue*

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed for the storage pool volume.

Detailed

Specifies that complete information is displayed for the storage pool volume.

Examples

Task 1

Display information on all storage pool volumes prefixed with the name ATF.

Command: Query Volume atf*

Output:

Volume Name	Storage Pool Name	Device Class Name	Estimated Capacity (MB)	%Util	Volume Status
ATF001	BACKUPPOOL	DISK	400.1	18.3	On-Line
ATF002	BACKUPPOOL	DISK	400.1	18.8	On-Line
ATF003	BACKUPPOOL	DISK	400.1	18.9	On-Line

where:

Volume Name

Specifies the name of the storage pool volume.

Storage Pool Name

Specifies the name of the storage pool to which the volume is defined.

Device Class Name

Specifies the name of the device class assigned to the storage pool.

Estimated Capacity (MB)

Specifies the estimated capacity of the volume, in megabytes.

For DISK devices, this value is the capacity of the volume.

For sequential access devices, this value is an estimate of the total space of the volume.

%Util

Specifies an estimate of the utilization of the volume.

For DISK devices, this value is specified as a percentage of the value for the estimated capacity, including cached data.

For sequential access devices, this value is specified as a percentage of the number of active bytes on the volume and the estimated capacity the volume.

Volume Status

Specifies whether the storage pool volume is available to ADSM.

Task 2

Display information in full detail about all storage pool volumes with names that begin with HARVEY.

Command: Query Volume harvey* Format=Detailed

Output:

```

        Volume Name: HARVEY.ADSM.STG1
        Storage Pool Name: BACKUPPOOL
        Device Class: DISK
        Estimated Capacity (MB): 5.8
            %Util: 0.1
        Volume Status: On-Line
        Access: Read/Write
        Pct. Reclaimable Space:
        Scratch Volume?:
        In Error State?:
        Number of Writable Sides:
        Number of Times Mounted:
        Write Pass Number:
        Approx. Date Last Written:
        Approx. Date Last Read:
        Date Became Pending:
        Number of Write Errors:
        Number of Read Errors:
        Volume Location:
        Last Update by (administrator): SERVER_CONSOLE
        Last Update Date/Time: 04/14/1995 13:32:26
    
```

```

        Volume Name: HARVEY.ADSM.STG2
        Storage Pool Name: BACKUPPOOL
        Device Class: DISK
        Estimated Capacity (MB): 11.7
            %Util: 0.0
        Volume Status: On-Line
        Access: Read/Write
        Pct. Reclaimable Space:
        Scratch Volume?:
        In Error State?:
        Number of Writable Sides:
        Number of Times Mounted:
        Write Pass Number:
        Approx. Date Last Written:
        Approx. Date Last Read:
        Date Became Pending:
        Number of Write Errors:
        Number of Read Errors:
        Volume Location:
        Last Update by (administrator): SERVER_CONSOLE
        Last Update Date/Time: 04/14/1995 13:32:29
    
```

where:

Volume Name

Specifies the name of the storage pool volume.

Storage Pool Name

Specifies the name of the storage pool to which the volume is defined.

Device Class Name

Specifies the name of the device class assigned to the storage pool.

Estimated Capacity (MB)

Specifies the estimated capacity of the volume, in megabytes.

For DISK devices, this value is the capacity of the volume.

For sequential access devices, this value is an estimate of the total space of the volume.

%Util

Specifies an estimate of the utilization of the volume.

For DISK devices, this value is specified as a percentage of the value for the estimated capacity, including cached data.

For sequential access devices, this value is specified as a percentage of the number of active bytes on the volume and the estimated capacity of the volume.

Volume Status

Specifies whether the volume is available to ADSM.

Access

Specifies whether the volume is available to ADSM.

Pct. Reclaimable Space

Specifies the amount of space on this volume that can be reclaimed. This value is compared to the value for the RECLAIM parameter (in DEFINE STGPOOL) to determine if reclamation is necessary.

Scratch Volume?

Specifies whether the volume is a scratch volume.

In Error State?

Specifies whether the volume is in an error state. Volumes in an error state cannot be written to.

Number of Writable Sides

This information is reserved for ADSM.

Number of Times Mounted

Specifies the number of times the volume has been mounted.

Write Pass Number

Specifies the number of times the volume has been written to from the beginning to the end of the sequential access medium.

Approx. Date Last Written

Specifies the approximate date on which the volume was last written.

Approx. Date Last Read

Specifies the approximate date on which the volume was last read.

Date Became Pending

Specifies the date that the volume's status changed to pending.

Number of Write Errors

Specifies the number of writing errors that have occurred on the volume.

Number of Read Errors

Specifies the number of reading errors that have occurred on the volume.

Volume Location

Specifies the location of the volume.

QUERY VOLUME

Last Update by (administrator)

Specifies the name of the administrator that has defined or most recently updated the volume.

Last Update Date/Time

Specifies when the volume was defined or most recently updated.

Related Commands

Table 107 provides references to related commands.

Table 107. Commands Related to QUERY VOLUME

Command	Description	Reference
DEFINE VOLUME	Assigns a volume to be used for storage within a specified storage pool.	Page 148
DELETE VOLUME	Deletes a volume from a storage pool.	Page 185
UPDATE VOLUME	Updates the attributes of one or more storage pool volumes.	Page 495
VARY	Specifies whether a disk volume is available to the server for use.	Page 504

QUIT (End the Interactive Mode of the Administrative Client)

Use the QUIT command to end an administrative client session in interactive mode.

You cannot use the QUIT command from the SERVER_CONSOLE administrative ID, or the console or batch modes of the administrative client.

For more information on the administrative client, see “Using the Administrative Client Session” on page 11.

Privilege Class

Any administrator can issue this command.

Syntax

```
▶▶—QUIT—◀◀
```

Parameters

None.

Examples

Task

End an administrative client session in the interactive mode.

Command: QUIT

Output:

```
ANS5103I Highest return code was n.
```

where *n* is the highest return code encountered during the administrative session.

Related Commands

None.

REDUCE DB (Decrease the Assigned Capacity of the Database)

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 4MB increments. If you do not specify the reduction in 4MB increments, ADSM rounds the number to the next 4MB partition. For example, if you reduce the database by 1MB, ADSM reduces the capacity of the database by 4MB.

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

The REDUCE DB command can be initiated while users are accessing the server. Use the QUERY DB command to determine how much you can reduce the assigned capacity of the database.

This command can be used to free up a database volume for deletion. You can specify the QUERY DBVOLUME command (FORMAT=DETAILED) to display the current utilization for your database volumes.

This command can generate a background process that can be canceled with the CANCEL PROCESS command. If a REDUCE DB background process is canceled, the database's assigned capacity may be partially reduced. To display information on background processes, use the QUERY PROCESS command.

Privilege Class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
►►—REDuce DB—nnn—◄◄
```

Parameters

nnn

Specifies the number of megabytes to be removed from the assigned capacity of the database. This parameter is required. This value is in increments of 4MB.

Examples

Task

Reduce the size of the database by 52MB.

Command: REDuce DB 52

Output:

```
ANR2250I Database assigned capacity has been reduced.
```

Related Commands

Table 108 provides references to related commands.

Table 108. Commands Related to REDUCE DB

Command	Description	Reference
CANCEL PROCESS	Cancels a background server process.	Page 64
DEFINE DBCOPY	Creates a volume copy of a database volume.	Page 95
DEFINE DBVOLUME	Assigns a volume to be used for database storage.	Page 97
EXTEND DB	Extends usable database storage within the volumes previously allocated for its use.	Page 213
QUERY DB	Displays allocation information about the database.	Page 274
QUERY DBVOLUME	Displays information about one or more database volumes.	Page 280
QUERY PROCESS	Displays information about one or all background server processes.	Page 332

REDUCE LOG (Decrease the Assigned Capacity of the Recovery Log)

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

The REDUCE LOG command can be initiated while users are accessing the server. Use the QUERY LOG command to determine how much you can reduce the assigned capacity of the recovery log.

The REDUCE LOG command can be used to free up a recovery log volume for deletion. You can use the QUERY LOGVOLUME command (FORMAT=DETAILED) to display the current utilization for your recovery log volumes.

This command can generate a background process that can be canceled with the CANCEL PROCESS command. If a REDUCE LOG background process is canceled, the recovery log's assigned capacity may be partially reduced. To display information on background processes, use the QUERY PROCESS command.

Privilege Class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
►►—REDuce LOG—nnn—◄◄
```

Parameters

nnn

Specifies the number of megabytes to be removed from the assigned capacity of the recovery log. This parameter is required. The value is in increments of 4MB.

Examples

Task

Reduce the size of the recovery log by 16MB.

Command: REDuce LOG 16

Output:

```
ANR2270I Recovery log assigned capacity has been reduced.
```


Related Commands

Table 109 provides references to related commands.

Table 109. Commands Related to REDUCE LOG

Command	Description	Reference
CANCEL PROCESS	Cancels a background server process.	Page 64
DEFINE DBBACKUPTRIGGER	Defines a trigger for automatic ADSM database backups.	Page 91
DEFINE LOGCOPY	Creates a volume copy of a recovery log volume.	Page 110
DEFINE LOGVOLUME	Defines a disk volume to the recovery log.	Page 112
EXTEND LOG	Extends usable log storage within the volumes previously allocated for its use.	Page 215
QUERY DBBACKUPTRIGGER	Displays information about the database backup trigger.	Page 278
QUERY LOG	Displays allocation information about the recovery log.	Page 307
QUERY LOGVOLUME	Displays information about one or more log volumes.	Page 311
QUERY PROCESS	Displays information about one or all background server processes.	Page 332
SET LOGMODE	Sets the mode for saving records in the recovery log.	Page 416

REGISTER ADMIN (Register an Administrator)

Use the REGISTER ADMIN command to add an administrator to the server. Once an administrator has been registered, the administrator can issue a limited set of commands, including all query commands. To provide additional administrator privileges, use the GRANT AUTHORITY command.

Privilege Class

To issue this command, you must have system privilege.

Syntax

```

▶▶—REGister Admin—adminname—password—[CONtact=contactinfo]—▶▶

```

Parameters

adminname

Specifies the name of the administrator to be registered. This parameter is required. The maximum length of the name is 64 characters.

password

Specifies the password of the administrator to be registered. This parameter is required. Use the SET PASSEXP command to specify how long an administrator or client node password is valid. The maximum length of the password is 64 characters.

CONtact=*contactinfo*

Specifies information identifying the administrator being registered. This parameter is optional. The maximum length of this string is 255 characters. The contact information must be enclosed in quotation marks if it contains any blanks.

Examples

Task

Define an administrator, LARRY, with the password PASSONE. You can identify LARRY as second-shift personnel by specifying this information with the CONTACT parameter.

Command: REGister Admin larry passone CONtact='second shift'

Output:

```
ANR2068I Administrator LARRY registered.
```

Related Commands

Table 110 provides references to related commands.

Table 110. Commands Related to REGISTER ADMIN

Command	Description	Reference
GRANT AUTHORITY	Assigns one or more privilege classes to an administrator.	Page 217
LOCK ADMIN	Prevents an administrator from accessing ADSM.	Page 244
QUERY ADMIN	Displays information about one or more ADSM administrators.	Page 256
REMOVE ADMIN	Removes an administrator from the list of registered administrators.	Page 377
RENAME ADMIN	Changes an ADSM administrator's name.	Page 380
SET PASSEXP	Specifies the number of days after which a password is expired and must be changed.	Page 423
UNLOCK ADMIN	Enables a locked administrator to access ADSM.	Page 435
UPDATE ADMIN	Changes the password or contact information associated with any administrator.	Page 437

REGISTER LICENSE (Register a New License)

Use the REGISTER LICENSE command to register a new license with the ADSM server, to change existing server license terms. You can register a new license for the following:


- Additional clients
- Environment support for desktop clients (clients other than UNIX and OpenEdition MVS clients; for example, DOS clients)
- Environment support for OpenEdition MVS clients
- Environment support for UNIX clients
- Support for space management
- Device support for specified removable media libraries and drives, including device support via the secondary server attachment feature

Note: The REGISTER LICENSE command adds a LICENSE statement to the end of the server options file if it completes successfully. To use this command, ensure that the options file as specified in the server startup JCL (OPT DD statement) is coded with a DISP=MOD parameter, and not DISP=SHR.

Privilege Class

To issue this command, you must have system privilege.

Syntax



Parameters

CLIENTS *number*

Specifies the licensed number of client connections for the server.

DESKTOP

Specifies that the server is licensed to support desktop clients (clients on platforms other than UNIX and OpenEdition MVS; for example, DOS clients).

OPENEDITION

Specifies that the server is licensed to support OpenEdition MVS clients.

UNIX

Specifies that the server is licensed to support UNIX clients.

SPACEMGMT

Specifies that the server is licensed to support space managed clients. If the server is not licensed for space management, migration operations for clients using space management fail.

DEVICEMOD3

Specifies that the server is licensed for Device Support Module 3.

DEVICEMOD4

Specifies that the server is licensed for Device Support Module 4.

SECONDARYSERVERATTACH

Specifies that the server is licensed for device support via a secondary server attachment feature.

Examples

Task

Register a new license with the ADSM server for support of 500 client connections.

Command:REGister LICense CLIENTS 500

Output:

```

ANR2852I Current license information:
ANR2835I Server is licensed for 25 clients.
ANR2844I Server is licensed to support DeskTop clients.
ANR2855I Server is licensed to support space-managed clients.
ANR2854I Server is licensed for device support module 2.
ANR2853I New license information:
ANR2835I Server is licensed for 500 clients.
ANR2844I Server is licensed to support DeskTop clients.
ANR2855I Server is licensed to support space-managed clients.
ANR2854I Server is licensed for device support module 2.

```

Related Commands

Table 111 provides references to related commands.

Table 111. Commands Related to REGISTER LICENSE

Command	Description	Reference
AUDIT LICENSES	Checks for compliance with defined licenses.	Page 43
QUERY LICENSE	Displays information about licenses and audits.	Page 304
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.	Page 342
SET LICENSEAUDITPERIOD	Specifies the number of days between automatic license audits.	Page 414

REGISTER NODE (Register a Client Node)

Use the REGISTER NODE command to register a client node to the server.

If your site has a closed registration system, then client nodes must be registered by an administrator with the REGISTER NODE command. You can use the REGISTER NODE command whether or not closed registration is in effect.

If your site is using open registration, client nodes can register themselves to the default policy domain named STANDARD, using only the default values for the attributes defined with the REGISTER NODE command.

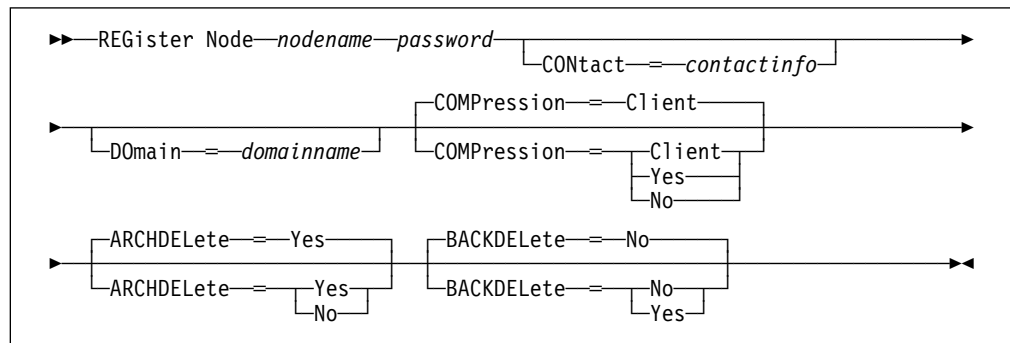
If a client requires a different policy domain than STANDARD, you must register the client node with this command.

The ARCHDELETE and BACKDELETE parameters determine whether the client node user can delete backup versions or archive copies from data storage. These parameters do not affect an administrator's ability to delete file spaces.

Privilege Class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the client node is assigned.

Syntax



Parameters

nodename

Specifies the name of the client node to be registered. This parameter is required. The maximum length of the name is 64 characters.

password

Specifies the client node's password. This parameter is required. The maximum length of the name is 64 characters. Use the SET PASSEXP command to specify how long an administrator or client node password is valid.

CONTACT=contactinfo

Specifies a text string of information identifying the node. This parameter is optional. The maximum length of the text string is 255 characters. The contact information must be enclosed in quotation marks if it contains any blanks.

DOMAIN=domainname

Specifies the name of the policy domain to which the node is assigned. This parameter is optional. If you do not specify a policy domain name, the node is assigned to the default policy domain (STANDARD).

COMPRESSION=compressionvalue

Specifies whether the client node compresses its files before sending them to the server for backup and archive. This parameter is optional. The default value is CLIENT. Possible values are:

Client

Specifies that the client's option file determines whether compression is used by the client node.

Yes

Specifies that the client node compresses its files before sending them to the server for backup and archive.

No

Specifies that the client node does not compress its files before sending them to the server for backup and archive.

ARCHDELETE=archdeletevalue

Specifies whether the client node can delete its own archived files from the server. This parameter is optional. The default value is YES. Possible values are:

Yes

Specifies that the client node can delete its own archive files from the server.

No

Specifies that the client node cannot delete its own archive files from the server.

BACKDELETE=backdeletevalue

Specifies whether the client node can delete its own backup files from the server. This parameter is optional. The default value is NO. Possible values are:

Yes

Specifies that the client node can delete its own backup files from the server.

No

Specifies that the client node cannot delete its own backup files from the server.

Examples

Task

Register the client node JOEOS2 with the password SECRETCODE and assign this node to the DOM1 policy domain. This node can delete its own backup and archive files from the server. All files are compressed by the client node before they are sent to the server.

Command: REGISTER Node joeos2 secretcode ARCHDELETE=Yes BACKDELETE=Yes
COMPRESSION=Yes

Output:

```
ANR2060I Node JOEOS2 registered.
```

Related Commands

Table 112 provides references to related commands.

Table 112. Commands Related to REGISTER NODE

Command	Description	Reference
DELETE FILESPACE	Deletes data associated with one or more client's filespaces.	Page 167
LOCK NODE	Prevents a client from accessing the server.	Page 245
REMOVE NODE	Removes a client from the list of registered nodes for a specific policy domain.	Page 378
RENAME NODE	Changes the name for an existing user that belongs to a specific policy domain.	Page 384
QUERY NODE	Displays partial or complete information about one or more clients.	Page 320
SET PASSEXP	Specifies the number of days after which a password is expired and must be changed.	Page 423
SET REGISTRATION	Specifies whether users can register themselves or must be registered by an administrator.	Page 428
UNLOCK NODE	Enables a locked user in a specific policy domain to access the server.	Page 436
UPDATE NODE	Changes the attributes associated with a client node.	Page 464

REMOVE ADMIN (Delete an Administrator)

Use the REMOVE ADMIN command to remove an administrator from the system.

You cannot remove the last system administrator from the system.

You cannot remove the SERVER_CONSOLE administrative ID.

Privilege Class

To issue this command you must have system privilege.

Syntax

```
▶▶—REMOve Admin—adminname————▶▶
```

Parameters

adminname

Specifies the name of the administrator to be removed. This parameter is required.

Examples

Task

Remove the administrator LARRY.

Command: REMOve Admin larry

Output:

```
Do you wish to proceed? (Y/N)
```

Type Y and press ENTER key. The following is displayed:

```
ANR2069I Administrator LARRY removed.
```

Related Commands

Table 113 provides references to related commands.

Table 113. Commands Related to REMOVE ADMIN

Command	Description	Reference
QUERY ADMIN	Displays information about one or more ADSM administrators.	Page 256
REGISTER ADMIN	Defines a new administrator without granting administrative authority.	Page 370

REMOVE NODE (Delete a Client Node)

Use the REMOVE NODE command to remove a client node from the server.

Before you can remove a client node, you must delete all backup and archived file spaces that belong to that client node. To remove a client node, you should first use the LOCK NODE command to ensure that the client node does not start a new session and begin backup and archive operations while you are deleting the client node's file spaces. Use the QUERY FILESPACE command to determine the file spaces for the client node. Then use DELETE FILESPACE command to delete all file spaces for the client node. After all file spaces that belong to the client node are deleted, you can remove the client node from ADSM.

Privilege Class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the client node is assigned.

Syntax

```
►►—REMOve Node—nodename—————►►
```

Parameters

nodename

Specifies the name of the client node to be removed. This parameter is required.

Examples

Task

Remove the client node LARRY.

Command: REMove Node larry

Output:

```
Do you wish to proceed? (Y/N)
```

Type Y and press the ENTER key. A message similar to the following is displayed:

```
ANR2061I Node LARRY removed from policy domain STANDARD.
```

Related Commands

Table 114 provides references to related commands.

Table 114. Commands Related to REMOVE NODE

Command	Description	Reference
DELETE FILESPACE	Deletes data associated with one or more client's filespaces.	Page 167
QUERY NODE	Displays partial or complete information about one or more clients.	Page 320
REGISTER NODE	Defines a client to the server and sets options for that user.	Page 374
RENAME NODE	Changes the name for an existing user that belongs to a specific policy domain.	Page 384

RENAME ADMIN (Rename an Administrator)

Use the RENAME ADMIN command to change the name of an existing administrator. Existing information for this administration such as password, contact information, and privilege classes is not be altered.

If you are assigning an existing administrative user ID to another person, it is recommended that you use the UPDATE ADMIN command to change the password.

This command is useful for assigning an existing administrative user ID to another person. This command can also be used when an employee changes his or her name and wants to be identified by a new user ID.

You cannot rename the SERVER_CONSOLE administrative ID.

Privilege Class

To issue this command you must have system privilege.

Syntax

```

  ►►—REName Admin—oldadminname—newadminname—————►►
  
```

Parameters

oldadminname

Specifies the name of the administrator to be renamed. This parameter is required.

newadminname

Specifies the new administrator name. This parameter is required. The maximum length of the name is 64 characters.

Examples

Task

Rename the ADSM administrator JOE to BILL.

Command: REName Admin joe bill

Output:

```

  ANR2070I Administrator JOE renamed to BILL.
  
```

Related Commands

Table 115 provides references to related commands.

Table 115. Commands Related to RENAME ADMIN

Command	Description	Reference
QUERY ADMIN	Displays information about one or more ADSM administrators.	Page 256
UPDATE ADMIN	Changes the password or contact information associated with any administrator.	Page 437

RENAME FILESPACE (Rename a Client Filespace on the Server)

Use the RENAME FILESPACE command to rename an existing client filespace on the server to a new filespace name.

This command is also useful after you have imported filespace into the server, before putting the filespace onto client nodes. Occasionally, the name of an imported filespace is the same as a filespace name that already exists on a client node. The server will not overlay an existing filespace that has the same name as an imported filespace. Instead, the imported filespace is given a new, system-generated filespace name. However, this new name may match a filespace name that already exists on the client node. (The client node may have filespace that have not yet been backed up and are unknown to the server.)

You can use the RENAME FILESPACE command to resolve this type of naming conflict or to rename the imported filespace to the naming convention used for the client node.

Privilege Class

Any administrator with unrestricted policy authority or with restricted policy authority over the client's policy domain can issue this command.

Syntax

```
►►—RENAmE FIlespacE—nodename—filespaceName—newfilespaceName—◄◄
```

Parameters

nodename

Specifies the name of the client node to which the filespace to be renamed belongs.

filespaceName

Specifies the name of the existing filespace defined for the client that is to be renamed.

A client filespace name is case sensitive; it must be specified using the same upper and lower case characters with which it is currently defined to the server.

newfilespaceName

Specifies the new name to be given to the client's filespace.

A client filespace name is case sensitive; it must be specified using the same upper and lower case characters with which it is to be defined to the server.

Usage Notes

1. Use the QUERY FILESPACE command to display the names of the filespace defined for each client in the server database. You can use this command to make sure that you are renaming a filespace that is defined for the client in the server database.
2. After importing filespace for a client node, this command can be used to rename the imported filespace to more appropriate names for the client, or to avoid any naming conflicts with existing client naming conventions.

Examples

Task

An AIX client node named MAX has backed up filespace /r033 to the ADSM server. The filespace is exported to tape, then imported to the server at a later time. When this filespace is imported, a system-generated name, /r031, is created for the imported filespace because /r033 already exists for client node MAX.

Client node MAX, however, already has a filespace named /r031 that has never been backed up and is unknown to the server. Unless the imported filespace is renamed, it will overlay filespace /r031 because the filespace name generated by the IMPORT function happened to be the same as a filespace on client node MAX that was unknown to the server.

The confusion that this IMPORT operation introduces for client node MAX is resolved by renaming imported filespace /r031 by using the following command:

Command: `RENAME FILESPACE MAX /r031 /imported-r033`

Output:

```
ANR0822I RENAME FILESPACE: Filespace r031 successfully
renamed imported-r033 for node MAX
```

The new name, /imported-r033, clearly identifies that the new filespace is an imported image of filespace /r033.

Related Commands

Table 116 provides references to related commands.

Table 116. Commands Related to RENAME FILESPACE

Command	Description	Reference
DELETE FILESPACE	Deletes data associated with one or more client's filespace.	Page 167
EXPORT NODE	Copies client node information to external media.	Page 199
QUERY FILESPACE	Displays information about data in one or more filespace that belong to a client.	Page 300
QUERY OCCUPANCY	Displays file space information by storage pool.	Page 324

RENAME NODE (Rename a Client Node)

Use the RENAME NODE command to rename a client node.

If you are assigning an existing client node ID to another person, it is recommended that you use the UPDATE NODE command to change the password.

This command is useful for assigning an existing client node user ID to another person. This command can also be used when an employee changes his or her name and wants to be identified by a new user ID.

Privilege Class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the client node is assigned.

Syntax

```
►►—REName Node—oldnodename—newnodename—◄◄
```

Parameters

oldnodename

Specifies the name of the client node to be renamed. This parameter is required.

newnodename

Specifies the new name of the client node. This parameter is required. The maximum length of the name is 64 characters.

Examples

Task

Rename the client node JOY to JOYCE.

Command: REName Node joy joyce

Output:

```
ANR2062I Node JOY renamed to JOYCE.
```


Related Commands

Table 117 provides references to related commands.

Table 117. Commands Related to RENAME NODE

Command	Description	Reference
QUERY NODE	Displays partial or complete information about one or more clients.	Page 320
UPDATE NODE	Changes the attributes associated with a client node.	Page 464

RESET BUFPOOL (Reset the Buffer Pool Statistics for the Database)

Use the RESET BUFPOOL command to reset the buffer pool statistics for the database. ADSM does not issue a response when you enter this command. The buffer pool statistics consist of:

- The cumulative number of database requests since the last RESET BUFPOOL command (or since the server was last started)
- The percentage of database requests for database pages that are already in the database buffer pool (as cached copies)
- The percentage of database requests that had to wait for a page in the database buffer pool to become available.

You can use the QUERY DB command (FORMAT=DETAILED) to display the current values for the buffer pool statistics.

Privilege Class

To issue this command, you must have system privilege or analyst privilege.

Syntax

```
►►—RESet BUFPool—————►►
```

Parameters

None.

Examples

Task

Reset the buffer pool statistics.

Command: RESet BUFPool

Output:

```
(no response)
```

Related Commands

Table 118 provides references to related commands.

Table 118. Commands Related to RESET BUFPOOL

Command	Description	Reference
QUERY DB	Displays allocation information about the database.	Page 274
RESET DBMAXUTILIZATION	Resets the maximum utilization statistic for the database.	Page 388
RESET LOGMAXUTILIZATION	Resets the maximum utilization statistic for the recovery log.	Page 390

RESET DBMAXUTILIZATION (Reset the Utilization Statistic for the Database)

Use the RESET DBMAXUTILIZATION command to reset the maximum utilization statistic for the database. ADSM does not issue a response when you enter this command.

This command resets the value of the maximum utilization statistic for the database to the current utilization. You can display the current value of the maximum utilization statistic for the database with the QUERY DB command.

Privilege Class

To issue this command, you must have system privilege or analyst privilege.

Syntax

```
▶▶—RESet DBMaxutilization—————▶▶
```

Parameters

None.

Examples

Task

Reset the maximum utilization statistic for the database.

Command: RESet DBMaxutilization

Output:

```
(no response)
```

Related Commands

Table 119 provides references to related commands.

Table 119. Commands Related to RESET DBMAXUTILIZATION

Command	Description	Reference
QUERY DB	Displays allocation information about the database.	Page 274
RESET BUFPOOL	Resets the buffer pool statistics for the database.	Page 386

RESET LOGCONSUMPTION (Reset the Cumulative Log Consumption Statistic)

Use the RESET LOGCONSUMPTION command to reset the statistic on the amount of recovery log space that has been consumed since the last reset.

You can run this command at any time.

This command allows an administrator to reset to zero the counter of megabytes of recovery log that have been consumed since the last reset. The recovery log consumption statistic is a cumulative measure of the amount of log space used by the server.

Privilege Class

To issue this command, you must have analyst privilege.

Syntax

```
▶▶—RESet LOGConsumption————▶▶
```

Parameters

None.

Examples

Task

Reset the statistic for recovery log consumption.

Command: RESet LOGConsumption

Output:

```
(no response)
```

Related Commands

Table 120 provides references to related commands.

Table 120. Commands Related to RESET LOGCONSUMPTION

Command	Description	Reference
QUERY LOG	Displays allocation information about the recovery log.	Page 307
SET LOGMODE	Sets the mode for saving records in the recovery log.	Page 416

RESET LOGMAXUTILIZATION (Reset the Utilization Statistic for the Log)

Use the RESET LOGMAXUTILIZATION command to reset the maximum utilization statistic for the recovery log. ADSM does not issue a response when you enter this command.

This command resets the value for the maximum utilization statistic of the recovery log to 0. You can display the current value of the maximum utilization statistic for the recovery log with the QUERY LOG command.

Privilege Class

To issue this command, you must have system privilege or analyst privilege.

Syntax

```
▶▶—RESet LOGMaxutilization—————▶▶
```

Parameters

None.

Examples

Task

Reset the maximum utilization statistic for the recovery log.

Command: RESet LOGMaxutilization

Output:

```
(no response)
```

Related Commands

Table 121 provides a reference to a related command.

Table 121. Command Related to RESET LOGMAXUTILIZATION

Command	Description	Reference
QUERY LOG	Displays allocation information about the recovery log.	Page 307

RESTORE STGPOOL (Restore Storage Pool Data from a Copy Pool)

Use the RESTORE STGPOOL command to restore files from one or more copy storage pools to a primary storage pool. The files must have been copied to the copy storage pools by using the BACKUP STGPOOL command.

When you use this command, ADSM restores all files that belong to the specified primary storage pool and that meet one of the following criteria:

- The primary copy of the file has been identified as having data-integrity errors during a previous operation.
- The primary copy of the file resides on a volume that has an access mode of DESTROYED.

After the files are restored, the old references to these files in the primary storage pool are deleted from the database. ADSM now locates these files on the volumes to which they were restored, rather than on the volumes on which they were previously stored. If a destroyed volume becomes empty because all files have been restored to other locations, the destroyed volume is automatically deleted from the database.

Note: Cached copies of files are never restored. Any cached files that have data-integrity errors or that reside on a *destroyed* volume are removed from the database during restore processing.

The restoration may be incomplete for one or more of the following reasons:

- Files were either never backed up or the backup copies are marked as damaged. Use the QUERY CONTENT command to get more information on the remaining files on the volume.
- A copy storage pool was specified on the RESTORE command, but files were backed up to a different copy storage pool. Use the PREVIEW parameter when you issue the RESTORE command again to determine if this is the problem.
- Volumes in the copy storage pool needed to perform the restore operation are offsite or unavailable. Check the activity log for messages that occurred during restore processing.
- Backup file copies in copy storage pools were moved or deleted by other ADSM processes during restore processing. To prevent this problem, do not issue the following commands for copy storage pool volumes while restore processing is in progress:
 - MOVE DATA
 - DELETE VOLUME (DISCARDATA=YES)
 - AUDIT VOLUME (FIX=YES)

Also, you can prevent reclamation processing for your copy storage pools by setting the RECLAIM parameter to 100 with the UPDATE STGPOOL command.

This command can be used to identify volumes that contain damaged, primary files. During restore processing, a message is issued for every volume in the restored storage pool that contains damaged, non-cached files. Use the QUERY CONTENT command to identify damaged, primary files on a specific volume.

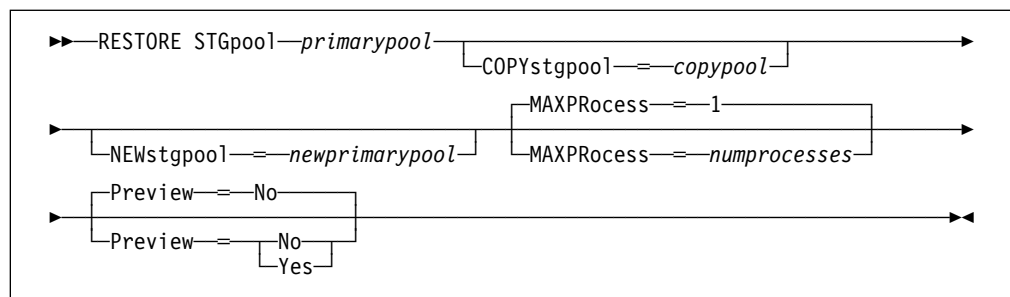
RESTORE STGPOOL

This command generates a background process that can be canceled with the CANCEL PROCESS command. If a RESTORE STGPOOL background process is canceled, some files may have already been restored prior to the cancellation. To display information on background processes, use the QUERY PROCESS command.

Privilege Class

To issue this command, you must have system privilege, unrestricted storage privilege, or restricted storage privilege for the primary storage pool for which files are to be restored. If you are a restricted storage administrator and you want to restore files to a new primary storage pool, you must also have authority for the new storage pool.

Syntax



Parameters

primarypool

Specifies the name of the primary storage pool that is being restored. This parameter is required.

COPYstgpool=*copypool*

Specifies the name of the copy storage pool from which the files are to be restored. This parameter is optional. If this parameter is not specified, files are restored from any copy pool in which copies can be located.

NEWstgpool=*newprimarypool*

Specifies the name of the new storage pool to which to restore the files. This parameter is optional. If this parameter is not specified, files are restored to the original primary storage pool (the pool being restored).

MAXPRocess=*numprocesses*

Specifies the maximum number of parallel processes that are used for restoring files. By using multiple, parallel processes, you may be able to improve throughput for the restore operation. This parameter is optional. You can specify a value from 1–999, inclusive. The default value is 1.

When determining the number of processes to use, consider the number of mount points and drives that can be dedicated to this operation. To access a volume in a sequential access device class, ADSM uses a mount point (a logical drive), and, if the sequential device class is not FILE, a physical drive. The number of mount points and drives available depends on other ADSM and system activity, and on the mount limits of the device classes for the storage pools that are involved in the restore operation. The maximum number of mount points for a device class is specified with the mount limit

parameter, which is the number of logical drives that ADSM can use at one time to access volumes in that device class.

Each process needs a mount point for copy storage pool volumes, and, if the device class is not FILE, each process also needs a drive. If the storage pool to which you are restoring files is a sequential storage pool, each process needs an additional mount point for primary storage pool volumes, and, if the device class is not FILE, an additional drive.

For example, suppose you specify MAXPROCESS=3 to restore files to a primary sequential storage pool from a copy storage pool, both assigned to the same CARTRIDGE device class. Each of the three processes requires two mount points and two drives. For all three processes to run, the CARTRIDGE device class must have a mount limit of no less than 6, and there must be at least 6 mount points and 6 drives currently available.

If you specify PREVIEW=YES, only one process is used, regardless of the value specified for the MAXPROCESS parameter. No mount points or drives are required when you specify PREVIEW=YES.

Preview=*value*

Specifies whether you only want to preview the restore operation, without actually performing it. This parameter is optional. The default value is NO. Possible values are:

No

Specifies that you want to perform the restore operation.

Yes

Specifies that you only want to preview the restore operation, without actually restoring the data. This option can be used to identify offsite volumes that would be required to restore a storage pool. This option displays:

- A list of volumes in the specified primary storage pool that contain damaged, primary files.
- The number of files and the number of bytes that will be restored, assuming that the access mode of the required copy storage pool volumes is READWRITE or READONLY when the restore operation is performed.
- A list of copy storage pool volumes containing files that will be restored. These volumes must be mounted if you actually perform the restore operation.
- A list of any volumes containing files that cannot be restored.

Note: If you preview a restore operation to obtain a list of offsite copy pool volumes that must be mounted, it is recommended that you change the access mode of the identified copy pool volumes to UNAVAILABLE. This prevents reclamation and MOVE DATA processing for these volumes until they have been transported to the onsite location for use in restore processing.

Examples

Task

Restore files from copy storage pool, JENCPSTG to the primary storage pool, JENSTG, and store them in a new storage pool named JENSTG2.

Command: RESTORE STGpool jenstg COPYstgpool=jencpstg NEWstgpool=jenstg2

Output:

```

ANR2110I RESTORE STGPOOL started as process 4.
ANR5220I Deleting FILE JENV03 from storage pool JENSTG.
ANR1230I Restore of primary storage pool JENSTG started
as process 5.
ANR1254I Removable volume JENV21 is required for data
storage restore.
ANR1254I Removable volume JENV22 is required for data
storage restore.
ANR5216I FILE JENV21 is expected to be mounted (R/O).
ANR5216I FILE JENV22 is expected to be mounted (R/W).
ANR5410I FILE volume JENV21 mounted.
ANR5410I FILE volume JENV31 mounted.
ANR5209I Dismounting volume JENV21 (read-only access).
ANR5410I FILE volume JENV22 mounted.
ANR5411I End-of-volume reached for FILE volume JENV31.
ANR5208I Dismounting volume JENV31 (updated).
ANR5410I FILE volume JENV32 mounted.
ANR1254I Removable volume JENV22 is required for data
storage restore.
ANR1234I Restore process 5 ended for storage pool JENSTG.
ANR5209I Dismounting volume JENV22 (read-only access).
ANR5208I Dismounting volume JENV32 (updated).
ANR1238I Restore of primary storage pool JENSTG has ended.
Files Restored: 1,
Bytes Restored: 1613227, Unreadable Files: 0, Unreadable
Bytes: 0.
ANR2208I Volume JENV01 deleted from storage pool JENSTG.
history information.
ANR2208I Volume JENV02 deleted from storage pool JENSTG.
ANR5220I Deleting FILE JENV01 from Storage Pool JENSTG.
ANR5220I Deleting FILE JENV02 from Storage Pool JENSTG.
    
```

Related Commands

Table 122 provides references to related commands.

Table 122 (Page 1 of 2). Commands Related to RESTORE STGPOOL

Command	Description	Reference
BACKUP STGPOOL	Backs up a primary storage pool to a copy storage pool.	Page 56
CANCEL PROCESS	Cancels a background server process.	Page 64
QUERY CONTENT	Displays information about one or more files currently residing in a storage pool volume.	Page 265

Table 122 (Page 2 of 2). Commands Related to RESTORE STGPOOL

Command	Description	Reference
QUERY PROCESS	Displays information about one or all background server processes.	Page 332
RESTORE VOLUME	Restores files stored on specified volumes in a primary storage pool from copy storage pools.	Page 396
UPDATE STGPOOL	Changes the attributes of a storage pool.	Page 482

RESTORE VOLUME (Restore Primary Volume Data from a Copy Pool)

Use the RESTORE VOLUME command to restore all files that are currently stored on one or more volumes in the same primary storage pool, and that were previously backed up to copy storage pools by using the BACKUP STGPOOL command. The files are restored from one or more copy storage pools. You can use this command to restore files on volumes known to be damaged, for example, a broken tape cartridge.

If more than one volume is specified to be restored, this command attempts to minimize volume mounts for the copy storage pool. Therefore, to restore more than one volume in the same primary storage pool, issue this command once and specify a list of volumes to be restored.

Note: Cached copies of files are never restored. Any cached files that reside on a volume that is being restored are removed from the database during restore processing.

Before you can use this command to restore a random-access volume, you must use the VARY command to vary the volume offline.

After files are restored, the old references to these files in the primary storage pool are deleted from the database. ADSM will now locate these files on the volumes to which they were restored, rather than on the volume on which they were previously stored.

This command changes the access mode of the specified volumes to DESTROYED. When the restoration is complete (all files on the volume are restored to other locations), the destroyed volume is empty and is then automatically deleted from the database.

The restoration may be incomplete for one or more of the following reasons:

- Files were either never backed up or the backup copies are marked as damaged. Use the QUERY CONTENT command to get more information on the remaining files on the volume.
- A copy storage pool was specified on the RESTORE command, but files were backed up to a different copy storage pool. Use the PREVIEW parameter when you issue the RESTORE command again to determine if this is the problem.
- Volumes in the copy storage pool needed to perform the restore operation are offsite or unavailable. Check the activity log for messages that occurred during restore processing.
- Backup file copies in copy storage pools were moved or deleted by other ADSM processes during restore processing. To prevent this problem, do not issue the following commands for copy storage pool volumes while restore processing is in progress:
 - MOVE DATA
 - DELETE VOLUME (DISCARDATA=YES)
 - AUDIT VOLUME (FIX=YES)

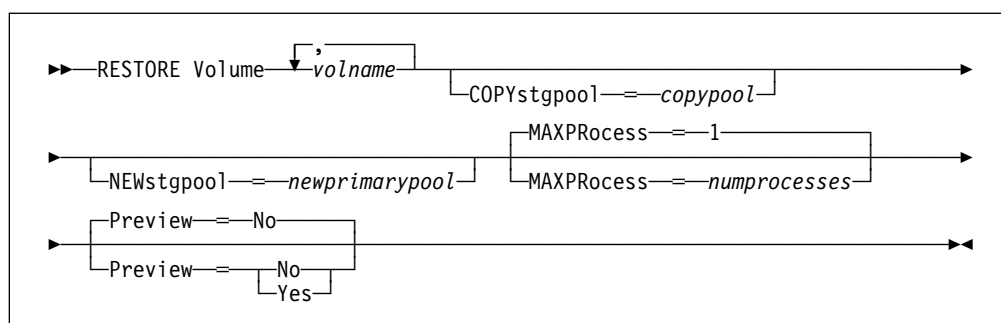
Also, you can prevent reclamation processing for your copy storage pools by setting the RECLAIM parameter to 100 with the UPDATE STGPOOL command.

This command generates a background process that can be canceled with the CANCEL PROCESS command. If a RESTORE VOLUME background process is canceled, some files may have already been restored prior to the cancellation. To display information on background processes, use the QUERY PROCESS command.

Privilege Class

To issue this command you must have system privilege, unrestricted storage privilege, or restricted storage privilege for the primary storage pool to which the specified volumes belong. If you are a restricted storage administrator and you want to restore files to a new primary storage pool, you must also have authority for the new storage pool.

Syntax



Parameters

volname

Specifies the name of the volume in a primary storage pool for which to restore files. This parameter is required. You can specify a list of volumes, if all volumes in the list belong to the same primary storage pool. The items in the list are separated by commas, with no intervening spaces.

COPYstgpool=*copypool*

Specifies the name of the copy storage pool from which the files are to be restored. This parameter is optional. If this parameter is not specified, files are restored from any copy pool in which copies can be located.

NEWstgpool=*newprimarypool*

Specifies the name of the new storage pool to which to restore the files. This parameter is optional. If this parameter is not specified, files are restored to the primary storage pool to which the volume or volumes being restored belong.

MAXProcess=*numprocesses*

Specifies the maximum number of parallel processes that are used for restoring files. By using multiple, parallel processes, you may be able to improve throughput for the restore operation. This parameter is optional. You can specify a value from 1–999, inclusive. The default value is 1.

When determining the number of processes to use, consider the number of mount points and drives that can be dedicated to this operation. To access a volume in a sequential access device class, ADSM uses a mount point (a logical drive), and, if the sequential device class is not FILE, a physical drive. The number of mount points and drives available depends on other ADSM

RESTORE VOLUME

and system activity, and on the mount limits of the device classes for the storage pools that are involved in the restore operation. The maximum number of mount points for a device class is specified with the mount limit parameter, which is the number of logical drives that ADSM can use at one time to access volumes in that device class.

Each process needs a mount point for copy storage pool volumes, and, if the device class is not FILE, each process also needs a drive. If the storage pool to which you are restoring files is a sequential storage pool, each process needs an additional mount point for primary storage pool volumes, and, if the device class is not FILE, an additional drive.

For example, suppose you specify MAXPROCESS=3 to restore files to a primary sequential storage pool from a copy storage pool, both assigned to the same CARTRIDGE device class. Each of the three processes requires two mount points and two drives. For all three processes to run, the CARTRIDGE device class must have a mount limit of no less than 6, and there must be at least 6 mount points and 6 drives currently available.

If you specify PREVIEW=YES, only one process is used, regardless of the value specified for the MAXPROCESS parameter. No mount points or drives are required when you specify PREVIEW=YES.

Preview=*value*

Specifies whether you want to preview the restore operation without actually restoring data. This parameter is optional. The default value is NO. Possible values are:

No

Specifies that you want to perform the restore operation.

Yes

Specifies that you want to preview the restore operation without actually restoring the data. This option displays:

- The number of files and the number of bytes that will be restored, assuming that the access mode of the required copy storage pool volumes is READWRITE or READONLY when the restore operation is performed.
- A list of copy storage pool volumes containing files that will be restored. These are volumes that must be mounted if you actually perform the restore.
- A list of any volumes containing files that are not restorable.

This option can be used to identify offsite volumes that would be required to restore a storage pool.

Note: If you preview this restore operation to obtain a list of offsite copy pool volumes that must be mounted, it is recommended that you change the access mode of the identified copy pool volumes to UNAVAILABLE. This prevents reclamation and MOVE DATA processing for these volumes until they have been transported to the onsite location for use in restore processing.

Examples

Task

Restore files stored on volumes, JENV31 and JENV32 belonging to the same primary storage pool, JENSTG2.

Command: RESTORE Volume jenv31,jenv32

Output:

```

ANR1232I Restore of volumes in primary storage pool JENSTG2
started as process 6.
ANR1254I Removable volume JENV21 is required for data storage restore.
ANR1254I Removable volume JENV22 is required for data storage restore.
ANR2114I RESTORE VOLUME: Access mode for volume JENV31 updated to
"destroyed".
ANR2114I RESTORE VOLUME: Access mode for volume JENV32 updated to
"destroyed".
ANR5216I FILE JENV21 is expected to be mounted (R/O).
ANR5216I FILE JENV22 is expected to be mounted (R/W).
adsm>
ANR5410I FILE volume JENV21 mounted.
ANR5410I FILE volume JENV33 mounted.
ANR5410I FILE volume JENV22 mounted.
ANR5411I End-of-volume reached for FILE volume JENV33.
ANR5208I Dismounting volume JENV33 (updated).
ANR5410I FILE volume JENV34 mounted.
ANR1254I Removable volume JENV22 is required for data storage restore.
ANR1235I Restore process 6 ended for volumes in storage pool
storage pool JENSTG2.
ANR5209I Dismounting volume JENV22 (read-only access).
ANR5208I Dismounting volume JENV34 (updated).
ANR1240I Restore of volumes in primary storage pool JENSTG2
has ended.
Files Restored: 1, Bytes Restored: 1613227, Unreadable
Files: 0, Unreadable Bytes: 0.
ANR2208I Volume JENV31 deleted from storage pool JENSTG2.
ANR2208I Volume JENV32 deleted from storage pool JENSTG2.
ANR5220I Deleting FILE JENV31 from storage pool JENSTG2.
ANR5220I Deleting FILE JENV32 from Storage pool JENSTG2.

```

RESTORE VOLUME

Related Commands

Table 123 provides references to related commands.

Table 123. Commands Related to RESTORE VOLUME

Command	Description	Reference
BACKUP STGPOOL	Backs up a primary storage pool to a copy storage pool.	Page 56
CANCEL PROCESS	Cancels a background server process.	Page 64
QUERY PROCESS	Displays information about one or all background server processes.	Page 332
RESTORE STGPOOL	Restores files to a primary storage pool from copy storage pools.	Page 391

REVOKE AUTHORITY (Remove Administrator Authority)

Use the REVOKE AUTHORITY command to revoke one or more privilege classes from an administrator. You can also use this command to reduce the number of policy domains to which a restricted policy administrator has authority. This command can also be used to reduce the number of storage pools to which a restricted storage administrator has authority.

At least one administrator must have system privilege; therefore, if the administrator is the only one with system privilege, you cannot revoke their authority.

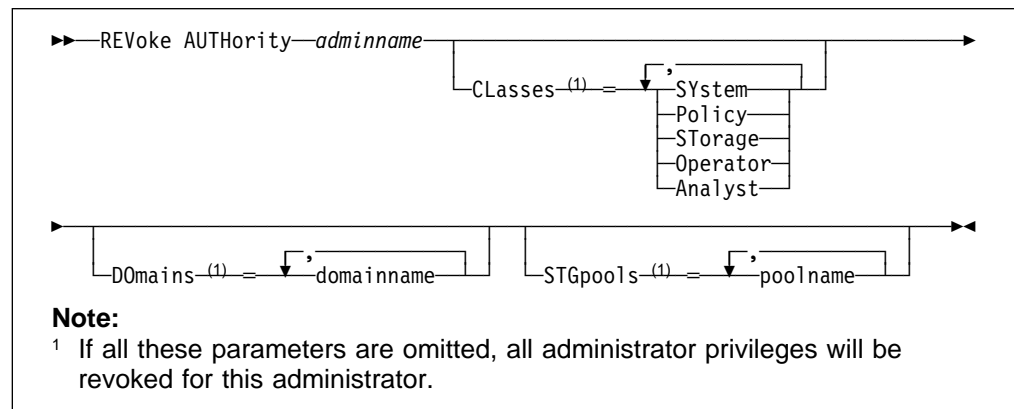
To change an unrestricted policy administrator to a restricted policy administrator, you must first use this command to revoke the unrestricted privilege. Then, use the GRANT AUTHORITY command to grant the administrator restricted policy privilege and to identify the policy domains to which the administrator has authority.

To change an unrestricted storage administrator to a restricted storage administrator, you must first use this command to revoke the unrestricted privilege. Then, use the GRANT AUTHORITY command to grant the administrator restricted storage privilege and to identify the storage pools to which the administrator has authority.

Privilege Class

To issue this command you must have system privilege.

Syntax



Parameters

adminname

Specifies the name of the administrator whose administrative privilege is to be revoked. This parameter is required.

Classes=classes

Specifies one or more administrative privilege classes to be revoked. You can specify more than one class by delimiting each with a comma.

System

Indicates that system authority is to be revoked for this administrator. If CLASSES=SYSTEM is specified, no other classes can be specified, and the DOMAINS and STGPOOLS parameters cannot be specified.

REVOKE AUTHORITY

Policy

Indicates that policy privilege is to be revoked for this administrator. To revoke all policy privilege, specify `CLASSES=POLICY` and do not specify the `DOMAINS` parameter.

STorage

Indicates that storage privilege is to be revoked for this administrator. To revoke all storage privilege, specify `CLASSES=STORAGE` and do not specify the `STGPOOLS` parameter.

Operator

Indicates that operator privilege is to be revoked for this administrator.

Analyst

Indicates that analyst privilege is to be revoked for this administrator.

DOMains=*domainlist*

Specifies a list of policy domains that can no longer be managed by a restricted policy administrator. (The administrator had been authorized to manage these domains until the `REVOKE` command was issued.) This parameter is optional. The items in the list are separated by commas, with no intervening spaces. You can use a pattern-matching expression to specify a name. Authority for all matching domains will be revoked. If `DOMAINS` is specified, the parameter `CLASSES=POLICY` is optional.

STGpools=*poollist*

Specifies a list of storage pools that can no longer be managed by a restricted policy administrator. (The administrator had been authorized to manage these storage pools until the `REVOKE` command was issued.) This parameter is optional. The items in the list are separated by commas, with no intervening spaces. You can use a pattern-matching expression to specify a name. Authority for all matching storage pools will be revoked. If `STGPOOLS` is specified then the parameter `CLASSES=STORAGE` is optional.

Usage Notes

- If you use the `REVOKE AUTHORITY` command without the `CLASSES`, `DOMAINS`, and `STGPOOLS` parameters, you will revoke all privileges for the specified administrator.
- You can revoke unrestricted policy privilege from an administrator by specifying the `CLASSES=POLICY` parameter. You cannot use the `DOMAINS` parameter to revoke authority for selected domains from an unrestricted policy administrator.
- You can revoke unrestricted storage privilege from an administrator by specifying the `CLASSES=STORAGE` parameter. You cannot use the `STGPOOLS` parameter to revoke authority for selected storage pools from an unrestricted storage administrator.

Examples

Task 1

Revoke part of administrator FRED's privileges. FRED has restricted policy privilege for the policy domains STANDARD, AIXDOMAIN1, and AIXDOMAIN2. Restrict FRED's policy privilege to the STANDARD policy domain.

Command: REVOke AUTHority FRED CLasses=Policy D0mains=AIX*

Output:

```
ANR2085I Restricted policy privilege revoked for
administrator FRED - policy domain AIXDOMAIN1.
ANR2085I Restricted policy privilege revoked for
administrator FRED - policy domain AIXDOMAIN2.
```

Task 2

Administrator SALLY currently has operator, analyst, and restricted policy privilege. Revoke all administrative privileges for administrator SALLY. To revoke all administrative privileges for an administrator, identify the administrator, but do not specify CLASSES, DOMAINS, or STGPOOLS. SALLY remains an administrator but she can only use those commands that can be issued by any administrator.

Command: REVOke AUTHority SALLY

Output:

```
ANR2085I Restricted policy privilege revoked for
administrator SALLY - policy domain AIXDOMAIN1.
ANR2088I Analyst privilege revoked for administrator SALLY.
ANR2089I Operator privilege revoked for administrator SALLY.
```

Task 3

Revoke system privilege for administrator PAULA. PAULA remains an administrator but she can only use those commands that can be issued by any administrator. (You can, if needed, use the GRANT AUTHORITY command to give privileges to PAULA.)

Command: REVOke AUTHority PAULA CLasses=System

Output:

```
ANR2083I System privilege revoked for administrator PAULA.
```

Task 4

Administrator JANE currently has operator, analyst, and unrestricted storage privilege. Revoke her operator and analyst privileges. After the command is issued, the administrator JANE continues to have unrestricted storage privilege.

REVOKE AUTHORITY

Command: REvOke AUTHority JANE Classes=Operator,Analyst

Output:

```
ANR2088I Analyst privilege revoked for administrator JANE.  
ANR2089I Operator privilege revoked for administrator JANE.
```

Related Commands

Table 124 provides references to related commands.

Table 124. Commands Related to REVOKE AUTHORITY

Command	Description	Reference
GRANT AUTHORITY	Assigns one or more privilege classes to an administrator.	Page 217
QUERY ADMIN	Displays information about one or more ADSM administrators.	Page 256

ROLLBACK (Rollback Uncommitted Changes in a Macro)

Use the ROLLBACK command within a macro to undo any processing changes made by commands that have been executed by the server but not yet committed to the database. After a change has been committed, it is permanent; it cannot be rolled back. This command is useful for testing macros.

Make sure that your administrative client session is not running with the ITEMCOMMIT option when using this command. For more information on controlling command processing in a macro, see “Controlling Command Processing in a Macro” on page 20.

Privilege Class

Any administrator can issue this command.

Syntax

```
▶▶—ROLLBACK—▶▶
```

Parameters

None.

Examples

Task

Run the REGN macro with the ROLLBACK command to verify that the macro works without committing any changes.

Macro Contents:

```
/* Macro to register policy administrators
 & grant authority */
REGister Admin sara hobby
GRant AUTHority sara CLasses=Policy
REGister Admin ken plane
GRant AUTHority ken CLasses=Policy
ROLLBACK      /* prevents any changes from being committed */
```

Command: MACRO regn

ROLLBACK

Output:

```
ANS5101I Server command: 'REGister Admin sara hobby'  
ANR2068I Administrator SARA registered.  
  
ANS5101I Server command: 'GRant AUTHority sara  
Classes=policy'  
ANR2077I Unrestricted policy privilege granted to  
administrator SARA.  
  
ANS5101I Server command: 'REGister Admin ken plane'  
ANR2068I Administrator KEN registered.  
  
ANS5101I Server command: 'GRant AUTHority ken Classes=policy'  
ANR2077I Unrestricted policy privilege granted to  
administrator KEN.  
  
ANS5101I Server command: 'ROLLBACK'
```

Related Commands

Table 125 provides references to related commands.

Table 125. Commands Related to ROLLBACK

Command	Description	Reference
COMMIT	Makes changes to the database permanent.	Page 68
MACRO	Runs a specified macro file.	Page 246

SET ACCOUNTING (Set Accounting Records On or Off)

Use the SET ACCOUNTING command to create an accounting record every time a client node session ends.

An accounting record tracks the amount of storage used by a client node session. See *ADSM Administrator's Guide* for more information on accounting.

Use the QUERY STATUS command to determine whether accounting records are generated. At installation, this value is set to OFF.

Privilege Class

To issue this command you must have system privilege.

Syntax

```
▶▶ Set ACCounting [ON | OFF] ▶▶
```

Parameters

ON

Specifies that the server creates an accounting record every time a client node session ends. You must specify either ON or OFF.

OFF

Specifies that the server does not create accounting records. You must specify either ON or OFF.

Examples

Task

Create an accounting record at the end of each client node session.

Command: Set ACCounting ON

Output:

```
ANR2091I Accounting parameter set to ON.
```

Related Commands

Table 126 provides references to related commands.

Table 126. Commands Related to SET ACCOUNTING

Command	Description	Reference
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.	Page 342

SET ACTLOGRETENTION (Set the Retention Period for the Activity Log)

Use the SET ACTLOGRETENTION command to specify the retention period, in days, for server activity log records.

This command allows you to define how long activity log information is retained in the database. Activity log information includes messages, such as:

- When client sessions start and end
- When storage pool migration starts and ends
- Diagnostic error messages

Because the activity log is stored in the database, the size of the activity log should be factored into the amount of space allocated for the database. The activity log size is dependent upon the length of time the activity log information is retained. When the length of time is increased, the amount of accumulated data also increases, so that additional database storage is required. When sufficient space is not available for activity log records, ADSM suspends recording, and sends messages about this condition to the server console. If you increase the size of the database, ADSM starts activity log recording again.

After a message has passed the specified retention period, the server automatically deletes it.

Use the QUERY ACTLOG command to search for messages in the activity log.

Use the QUERY STATUS command to display the current retention period for the activity log. At installation, this value is set to 1 day.

Privilege Class

To issue this command you must have system privilege.

Syntax

```
►►—Set ACTlogretention—days—————◄◄
```

Parameters

days

Specifies an integer that defines the number of days to retain messages in the activity log. This parameter is required. You can specify an integer in the range of 0 to 9999, inclusive. A value of 0 indicates that activity log records will be retained only for the current day.

Examples

Task

Set the server to retain activity log records for 30 days.

Command: Set ACTlogretention 30

SET ACTLOGRETENTION

Output:

```
ANR2090I Activity log retention period set to 30 day(s).
```

Related Commands

Table 127 provides a reference to a related command.

Table 127. Command Related to SET ACTLOGRETENTION

Command	Description	Reference
QUERY ACTLOG	Searches for and displays messages from the server activity log.	Page 252

SET AUTHENTICATION (Set Password Authentication)

Use the SET AUTHENTICATION command to specify whether administrators and client nodes must use a password to access the server.

Use the QUERY STATUS command to display the current authentication setting. At installation, this value is set to ON.

Privilege Class

To issue this command you must have system privilege.

Syntax

```

▶▶—Set Authentication—ON
                        OFF
▶▶

```

Parameters

ON

Specifies that administrators and client nodes need a password to access the server.

OFF

Specifies that administrators and client nodes do not need a password to access the server.

Examples

Task

Set authentication to ON to require administrators and client nodes to enter a password when accessing the server.

Command: Set Authentication ON

Output:

```
ANR2095I Authentication parameter set to ON.
```

Related Commands

Table 128 provides a reference to a related command.

Table 128. Command Related to SET AUTHENTICATION

Command	Description	Reference
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.	Page 342

SET EVENTRETENTION (Specify the Retention Period for Event Records)

Use the SET EVENTRETENTION command to specify the retention period for event records in the server database. An event record is created whenever processing of a scheduled command is started or missed.

Event records are automatically removed from the database once both of the following conditions have been met:

- The specified retention period has passed from the scheduled execution date and time for the event
- The startup window for the event has elapsed.

Use this command to set a retention period for event records that will permit you to effectively monitor completed schedules. You can adjust the length of time event information is maintained to avoid insufficient or outdated data. To display event records, use the QUERY EVENT command.

You can use the DELETE EVENT command to delete event records regardless of whether their retention period has elapsed.

You can display the value for the event retention period with the QUERY STATUS command. At installation, this value is set to 10 days.

Privilege Class

To issue this command you must have system privilege.

Syntax

```
►►—Set EVentretention—days—◄◄
```

Parameters

days

Specifies the number of days for which event records are retained in the database. This parameter is required. You can specify an integer in the range of 0 to 9999, inclusive. A value of 0 indicates that only event records for the current day will be retained.

Examples

Task

Maintain event records in the database for a period no longer than 15 days.

Command: Set EVentretention 15

Output:

```
ANR2521I Event record retention period set to 15 day(s).
```

Related Commands

Table 129 provides references to related commands.

Table 129. Commands Related to SET EVENTRETENTION

Command	Description	Reference
DELETE EVENT	Deletes event records prior to a specified date and time.	Page 164
QUERY EVENT	Displays information about scheduled and completed events for selected clients.	Page 290
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.	Page 342


SET LICENSEAUDITPERIOD (Set License Audit Period)

Use the SET LICENSEAUDITPERIOD command to specify the period, in days, between automatic license audits performed by the ADSM server.

Privilege Class

To issue this command, you must have system privilege.

Syntax



Parameters

days

Specifies the number of days between automatic server license audits on the ADSM server. This parameter is optional. The default value is 30. You can specify an integer in the range of 1 to 30, inclusive.

Examples

Task

Specify the license audit period for every 14 days.

Command: Set LICenseauditperiod 14

Output:

```
ANR2814I SET LICENSEAUDITPERIOD: License Audit Period
changed to 14 days.
```

Related Commands

Table 130 provides references to related commands.

Table 130. Commands Related to SET LICENSEAUDITPERIOD

Command	Description	Reference
AUDIT LICENSES	Checks for compliance with defined licenses.	Page 43
QUERY AUDITOCCUPANCY	Displays the data storage utilization for a client node.	Page 262
QUERY LICENSE	Displays information about licenses and audits.	Page 304
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.	Page 342
REGISTER LICENSE	Registers a new license with the ADSM server.	Page 372

SET LOGMODE (Set Mode for Saving Recovery Log Records)

Use the SET LOGMODE command to set the mode for saving recovery log records. The log mode you specify determines how long ADSM saves records in the recovery log, which determines what kind of database recovery you can use. If the log mode is set to NORMAL, you can only restore a database to a specific point in time (point-in-time recovery). If the log mode is set to ROLLFORWARD, you can choose to restore a database to a specific point in time or to its most current state through the use of recovery log records (roll-forward recovery).

If the log mode is set to NORMAL, ADSM saves only those records needed to restore the database to the point of the last backup and resume database processing, if a system failure occurs. 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.

If the log mode is set to 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 when they become unnecessary because the database is successfully backed up. The recovery log records can be used to restore a database to its most current state after loading the most current database backup series.

A database backup series created while the log mode is set to ROLLFORWARD can be used to perform point-in-time recovery or roll-forward recovery.

If you specify ROLLFORWARD for the log mode, you can use the DEFINE DBBACKUPTRIGGER command to specify how much space can be used in the recovery log before ADSM automatically runs a backup of the database and removes any unnecessary records from the recovery log. Setting a database backup trigger helps to ensure that the recovery log does not run out of space for new records before the next backup is run.


For more information on how the log mode affects database recovery, see *ADSM Administrator's Guide*.

You cannot use the SET LOGMODE command while the BACKUP DB command is running. You must wait for the backup to complete or cancel the background process started by the BACKUP DB command before you can change the log mode.

Privilege Class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax



Parameters

Normal

Specifies that ADSM saves only those recovery log records that are needed to recover the database in case of a system failure. This is the default for the log mode.

When the log mode is NORMAL, no automatic backups are performed based on the database backup trigger. ADSM deletes recovery log records that were required to perform roll-forward recovery or to restore a single database volume. However, you can still use existing database backups to perform point-in-time recovery. ADSM keeps log records necessary to resume database processing if a system failure occurs.

In response to the change to NORMAL mode, ADSM may issue a message warning that only point-in-time database recovery will be possible. If the database backup trigger is set, the message also warns that the trigger will be disabled. If issued from an administrative client (not from the server console), ADSM asks whether to proceed with the change to NORMAL mode.

Rollforward

Specifies that ADSM saves enough log records to ensure that the database can be restored to its most current state using the most recent database backup series and recovery log records. If issued from an administrative client (not from the server console), ADSM asks whether to proceed with the change to ROLLFORWARD mode.

If the database backup trigger is defined, ADSM automatically and immediately runs a full backup.

If the database backup trigger is not defined when you change the log mode from NORMAL to ROLLFORWARD, a full backup is required the next time a backup is run. If you do not specify TYPE=FULL the next time you issue the BACKUP DB command, the command fails.

Examples

Task 1

Change the log mode from roll-forward to normal after the database backup trigger has been defined.

Command: Set LOGMode Normal

Output:

```
ANR2291W Changing the log mode to NORMAL will disable the database
backup trigger.
ANR2290W Changing the log mode to NORMAL will prevent roll-forward
recovery.
Do you wish to proceed? (Yes/No)
```

Type Y and press the Enter key. The following message is displayed:

```
ANR2296I Log mode set to NORMAL; database backup trigger disabled.
ANR2293I Only point-in-time database recovery is now possible.
```

SET LOGMODE

Task 2

Change the log mode from normal to roll-forward. No database backup trigger is defined.

Command: Set LOGMode Rollforward

Output:

```
ANS5100I Session established with server ADSM: MVS
ANR2294I Log mode set to ROLLFORWARD.
```

Related Commands

Table 131 provides references to related commands.

Table 131. Commands Related to SET LOGMODE

Command	Description	Reference
BACKUP DB	Backs up the ADSM database to sequential access volumes.	Page 50
DEFINE DBBACKUPTRIGGER	Defines a trigger for automatic ADSM database backups.	Page 91
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.	Page 342
UPDATE DBBACKUPTRIGGER	Changes one or more attributes of the database backup trigger.	Page 447

SET MAXCMDRETRIES (Set Number of Times Scheduler Retries Commands)

Use the SET MAXCMDRETRIES command to specify the maximum number of times that a client scheduler retries after a failed attempt to execute a scheduled command.

The MAXCMDRETRIES parameter only restricts the number of attempts that are made to retry a scheduled command that fails. This parameter does not limit the number of attempts made by the client scheduler to contact the server to obtain a schedule.

The MAXCMDRETRIES parameter can be specified by each user at the time their client scheduler program is started. Refer to the appropriate client documentation for more information on setting the maximum command retries value from the client.

You can use the SET MAXCMDRETRIES command to set a global value for the maximum number of retries, which overrides the value specified by the user. The client's value is overridden only if the client can contact the server.

You can display the current value for the maximum command retries with the QUERY STATUS command. At installation, ADSM is configured so that each client determines its own value for this setting.

Privilege Class

To issue this command you must have system privilege.

Syntax

```

  >> Set MAXCMDRetries maxcmdretries <<
  
```

Parameters

maxcmdretries

Specifies the maximum number of times the client scheduler retries commands after a failed attempt to execute a scheduled command. This parameter is optional. You can specify an integer in the range of 0 to 9999, inclusive. If you do not specify a value for this parameter, each client determines its own value for this parameter.

Examples

Task 1

Have all client schedulers retry a failed attempt to execute a scheduled command only twice.

Command: Set MAXCMDRetries 2

Output:

```

  ANR2525I Maximum number of command retries set to 2.
  
```

SET MAXCMDRETRIES

Task 2

Have each client determine how often to retry a failed command.

Command: Set MAXCMDRetries

Output:

```
ANR2526I Maximum number of command retries reset to value
determined by each client.
```

Related Commands

Table 132 provides a reference to a related command.

Table 132. Command Related to SET MAXCMDRETRIES

Command	Description	Reference
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.	Page 342

SET MAXSCHEDESESSIONS (Set Maximum Scheduled Sessions)

Use the SET MAXSCHEDESESSIONS command to regulate the number of sessions that the server can use for processing scheduled work. This command specifies the maximum number of scheduled sessions as a percentage of the total number of server sessions available.

This setting should be changed if client nodes receive messages that server sessions are not available when trying to execute scheduled events. If you find that scheduled sessions are not contacting the server, you may have a network error.

If you find that the number of sessions for scheduled operations is insufficient, you can increase either the total number of sessions or the maximum number of scheduled sessions. However, increasing the total number of sessions available can adversely affect server performance, and increasing the maximum number of scheduled sessions can reduce the server's opportunity to execute nonscheduled work.

You can use the QUERY STATUS command to display the current value for the maximum number of sessions and the maximum number of scheduled sessions. At installation, the value for the maximum number of scheduled sessions is set to 50 percent.

Privilege Class

To issue this command you must have system privilege.

Syntax

```
►►—Set MAXSchedsessions—schedulepercent—————►◄
```

Parameters

schedulepercent

Specifies an integer that gives the percentage of total server sessions that can be used for the processing of scheduled work. This parameter is required. You can specify an integer in the range of 0 to 100, inclusive. The maximum number of total available server sessions is determined by the MAXSESSIONS parameter in the server's option file.

If you set the maximum number of scheduled sessions to a value of 0, no scheduled events can begin. If you set the maximum number of scheduled sessions to a value of 100, the MAXSESSIONS parameter determines the maximum number of scheduled sessions.

Examples

Task

By setting MAXSESSIONS to 80, you limit the server to 80 active sessions at one time. You want only 25 percent of your total available sessions to be available for scheduled activity. The maximum number of client-server sessions available for scheduled work at any given time is 20 sessions.

SET MAXSCHEDESESSIONS

Command: Set MAXSCHedsessions 25

Output:

```
ANR2520I Scheduled sessions set to 25 percent.
```

Related Commands

Table 133 provides references to related commands.

Table 133. Commands Related to QUERY STATUS

Command	Description	Reference
QUERY OPTION	Displays information about one or more options from the server options file.	Page 327
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.	Page 342

SET PASSEXP (Set Password Expiration Date)

Use the SET PASSEXP command to set the password expiration period. The password expiration period applies to both administrator and client node passwords.

Use the QUERY STATUS command to display the current password expiration period. At installation, this value is set to 90 days.

Privilege Class

To issue this command, you must have system privilege.

Syntax

```
▶▶—Set PASSExp—days————▶▶
```

Parameters

days

Specifies an integer that defines the number of days for which the password remains valid. This parameter is required. You can specify an integer in the range of 1 to 9999, inclusive. If an administrator or client node password expires, the server prompts for a new password when the administrator or client node attempts to contact the server.

Examples

Task

Set the administrator and client node password expiration period to 45 days.

Command: Set PASSExp 45

Output:

```
ANR2092I Password expiration period set to 45 day(s).
```

Related Commands

Table 134 provides references to related commands.

Table 134. Commands Related to SET PASSEXP

Command	Description	Reference
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.	Page 342

SET QUERYSCHEDPERIOD (Set Query Period for Polling Client Nodes)

Use the SET QUERYSCHEDPERIOD command to regulate the frequency with which client nodes contact the server to obtain scheduled work when they are running in the client-polling mode.

The value for the QUERYSCHEDPERIOD parameter can be set by each client node at the time the client scheduler program is started.

You can set a global value for the period between attempts by the client to contact the server for scheduled work. This value overrides the value specified by the client.

The client's value is only overridden if the client can contact the server.


You can set this period to correspond to the frequency with which the schedule or client association changes are being made. If client nodes poll more frequently for schedules, changes to scheduling information (through administrator commands) are propagated more quickly to client nodes. However, increased polling by the client nodes also increases network traffic.

You can display the value for the period between schedule queries with the QUERY STATUS command. At installation, ADSTM is configured so that each client node determines its own value for this setting.

Privilege Class

To issue this command you must have system privilege.

Syntax



Parameters

queryschedperiod

Specifies the maximum number of hours between attempts by the client to contact the server for scheduled commands. This parameter is optional. You can specify an integer in the range of 1 to 9999, inclusive. If you do not specify a value for this parameter, each client determines its own value for this parameter.

Examples

Task 1

Have all clients using the POLLING scheduling mode contact the server every 24 hours.

Command: Set QUERYSCHEdperiod 24

Output:

```
ANR2523I Schedule query period set to 24 hour(s).
```


Task 2

Have all clients using the POLLING scheduling mode set their own query scheduling period.

Command: Set QUERYSChedperiod

Output:

ANR2524I Schedule query period reset to value determined by client.

Related Commands

Table 135 provides references to related commands.

Table 135. Commands Related to SET QUERYSCHEDPERIOD

Command	Description	Reference
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.	Page 342
SET SCHEDMODES	Specifies the central scheduling mode for the server.	Page 432

SET RANDOMIZE (Randomly Distribute Scheduled Start Times)

Use the SET RANDOMIZE command to specify the degree to which start times are randomized within the startup window of each schedule for clients using the client-polling mode.

Each schedule has an associated window during which execution can begin. To balance network and server load, the start times for individual clients can be scattered across that window. You can use this command to specify the fraction of the window over which start times for individual clients are to be distributed. The randomization occurs at the beginning of the window to allow time for retries, if necessary, unless the client's first contact with the server is after the start time for the event.

Set the randomization percentage to a value greater than 0 to prevent communication errors that can result from a large group of clients contacting the server simultaneously. If you do experience communication errors, you can increase the randomization percentage so that client contact is spread out, thus decreasing the chance for communication overload and failure.

It is possible, especially after a client or the server has been restarted, that a client may not poll the server until *after* the beginning of the time window in which the next scheduled event is to start. In this case, the starting time is randomized over the specified percentage of the *remaining* duration of the start window.

For example, consider the following situation:

- Randomization is set at 50 percent
- The startup window for a particular event is from 8:00 to 9:00
- Ten clients are associated with the schedule
- Nine clients poll the server before 8:00
- One client does not poll the server until 8:30.

The nine clients that polled the server *before* the beginning of the startup window are assigned randomly selected starting times between 8:00 and 8:30. The client that polled at 8:30 receives randomly selected starting time that is between 8:30 and 8:45.

You can display the value for the schedule randomization percentage with the QUERY STATUS command. At installation, this value is set to 50 percent.

Privilege Class

To issue this command, you must have system privilege.

Syntax

```
►►—Set RANDomize—randomizepercent—————◄◄
```

Parameters

randomizepercent

Indicates the percentage of the startup window over which the start times for individual clients are distributed. This value is specified as an integer in the range of 0 to 50, inclusive.

A value of 0 indicates that no randomization occurs and that all clients execute the schedule at the beginning of the window. A value of 50 indicates that clients are assigned start times that are randomly scattered across the first half of the startup window.

At installation, this value is set to 25, indicating that the first 25 percent of the window is used for randomization.

If you have specified DURUNITS=INDEFINITE in the DEFINE SCHEDULE command, the percentage specified by the RANDOMIZEPERCENT parameter is regarded as a fraction of a one-day time period. That is, if the value for the RANDOMIZEPERCENT parameter is 25, a schedule with an indefinite window duration has start times scattered over a 6-hour time range (0.25×24).

Examples

Task

Set randomization to 50 percent.

Command: Set RANDomize 50

Output:

```
ANR2522I Randomization set to 50 percent.
```

Related Commands

Table 136 provides references to related commands.

Table 136. Commands Related to SET RANDOMIZE

Command	Description	Reference
DEFINE SCHEDULE	Defines a schedule for a client operation or an administrative command.	Page 121
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.	Page 342
SET SCHEDMODES	Specifies the central scheduling mode for the server.	Page 432

SET REGISTRATION (Set Open or Closed Registration)

Use the SET REGISTRATION command to either require a system or policy administrator to register client nodes (closed registration), or permit each user to register his or her own workstation as a client node with the server (open registration).

With *closed* registration, a system or policy administrator defines:

- Each workstation's node name
- Each user's password
- The policy domain to which the client node belongs
- Whether the user can choose to compress files before sending them to server storage
- Whether the user is allowed to delete backup or archive files from server storage

With *open* registration, when a user attempts to access the server from an unregistered client node, the server prompts the user for a node name and password, and registers the workstation. The server sets the following defaults:

- Each client node is assigned to the policy domain named STANDARD.
- Each user defines whether data compression is used before files are sent to server storage.
- Each user is allowed to delete archived files from server storage.


Existing registered client nodes are not affected by changes in the registration process.

Use the QUERY STATUS command to display the status of registration. At installation, this value is set to CLOSED.

Privilege Class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the client node belongs.

Syntax



Parameters

Closed

Specifies that client nodes must be registered through a system or policy administrator.

Open

Specifies that users can register their workstations as client nodes with the server.

Examples

Task 1

Have a policy or system administrator register all client nodes to the server.

Command: Set REGistration Closed

Output:

```
ANR2093I Registration parameter set to CLOSED.
```

Task 2

Have all users register themselves as client nodes to the server.

Command: Set REGistration Open

Output:

```
ANR2093I Registration parameter set to OPEN.
```

Related Commands

Table 137 provides a reference to a related command.

Table 137. Command Related to SET REGISTRATION

Command	Description	Reference
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.	Page 342

SET RETRYPERIOD (Set Time between Retry Attempts)

Use the SET RETRYPERIOD command to specify the number of minutes between attempts by the client scheduler to retry after a failed attempt to contact the server or to execute a scheduled command.

Each client can set their own retry period at the time their client scheduler program is started. You can use this command to set a global value for the retry period which will override the value specified by all clients. The client's value is overridden only if the client is able to connect with the server.

When setting the period between retry attempts, set a time period that permits more than one retry attempt within a typical startup window.


This command is used in conjunction with the SET MAXCMDRETRIES command to regulate the period of time and the number of retry attempts to execute a failed command.

You can display the value for the period between retries by issuing the QUERY STATUS command. At installation, ADSM is configured so that each client determines its own value for this setting.

Privilege Class

To issue this command, you must have system privilege.

Syntax



Parameters

retryperiod

Specifies the number of minutes between attempts by the client scheduler to retry after a failed attempt to contact the server or to execute a scheduled command. You can specify an integer in the range of 1 to 9999, inclusive. If you do not specify a value for this parameter, each client determines its own value for this parameter.

Examples

Task 1

Allow each client to specify their own period between retries.

Command: Set RETRYPeriod

Output:

```
ANR2528I Retry period set to value determined by each client.
```

Task 2

Have the client scheduler retry failed attempts to contact the server or to execute scheduled commands every 15 minutes.

Command: Set RETRYPeriod 15

Output:

```
ANR2527I Retry period set to 15 minute(s).
```

Related Commands

Table 138 provides references to related commands.

Table 138. Commands Related to SET RETRYPERIOD

Command	Description	Reference
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.	Page 342
SET MAXCMDRETRIES	Specifies the maximum number of retries after a failed attempt to execute a scheduled command.	Page 419

SET SCHEDMODES (Select a Central Scheduling Mode)

Use the SET SCHEDMODES command to select the central scheduling modes that the server allows. Use this command when you set up central scheduling to determine how the clients communicate with the server to begin scheduled work.

You can restrict clients from using a scheduling mode with this command.

With client-polling scheduling, a client queries the server at prescribed time intervals to obtain scheduled work or to ensure that the schedules they are waiting to execute have not changed. At the scheduled start time, the client performs the scheduled backup or archive operations, sending the results to the server.

With server-prompted scheduling, clients wait to be contacted by the server when scheduled work needs to be performed and a session is available. When contacted, the client performs the scheduled backup or archive operations, sending the results to the server. The server-prompted mode can only be used between clients and servers that communicate through TCP/IP.

To select both server-prompted and client-polling modes, specify ANY.


It is very important that you communicate to all users the scheduling modes allowed by the server.

You can display the value for the scheduling mode supported by issuing the QUERY STATUS command. At installation, this value is set to ANY.

Privilege Class

To issue this command, you must have system privilege.

Syntax



Parameters

ANY

Specifies that the server allows clients to run in either the client-polling or the server-prompted scheduling mode. Each client must be configured to select the scheduling mode in which it operates.

POLLing

Specifies that only the client-polling mode can be used. Client nodes poll the server at prescribed time intervals to obtain scheduled work. If this mode is selected, configure each client to use the client-polling mode.

PRompted

Specifies that only the server-prompted mode can be used. This mode is only available for clients that communicate with TCP/IP. Client nodes wait to be contacted by the server when scheduled work needs to be performed and a session is available. If this mode is selected, configure each client to use the server-prompted mode.

Examples

Task 1

A server allows both server-prompted and client-polling central scheduling. You want to temporarily restrict the scheduled operations to clients that use the client-polling mode. If you set the schedule mode to POLLING, the server discontinues prompting clients to execute scheduled commands. This means that any client scheduler using the server-prompted mode waits until you set the schedule mode to ANY or PROMPTED.

Command: Set SCHEDMODEs POLLing

Output:

```
ANR2529I Scheduling mode set to POLLING.
```

Task 2

Set the central scheduling mode to PROMPTED. Only clients that communicate with TCP/IP will perform scheduled operations.

Command: Set SCHEDMODEs PRompted

Output:

```
ANR2530I Scheduling mode set to PROMPTED.
```

Related Commands

Table 139 provides a reference to a related command.

Table 139. Command Related to SET SCHEDMODES

Command	Description	Reference
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.	Page 342

SET SERVERNAME (Specify the Server Name)

Use the SET SERVERNAME command to specify the name by which the server is to be identified. At installation, the name of the server is set to ADSM. After installation, this command can be used to modify the name of the server.

You can use the QUERY STATUS command to see the name of the server.

The server name is reported to a client node when it contacts the server.

Privilege Class

To issue this command you must have system privilege.

Syntax

```
►►—Set SERVername—servername—————◄◄
```

Parameters

servername

Specifies the name by which the server is to be identified. This parameter is required. The maximum length of the name is 64 characters.

Examples

Task

Name the server WELLS_DESIGN_DEPT.

Command: Set SERVername wells_design_dept

Output:

```
ANR2094I Server name set to WELLS_DESIGN_DEPT.
```

Related Commands

Table 140 provides a reference to a related command.

Table 140. Command Related to SET SERVERNAME

Command	Description	Reference
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.	Page 342

UNLOCK ADMIN (Unlock an Administrator)

Use the UNLOCK ADMIN command to allow a locked administrator to re-access the server.

By using the LOCK ADMIN and UNLOCK ADMIN commands, you can temporarily restrict an administrator from accessing the server.

Privilege Class

To issue this command you must have system privilege.

Syntax

```
▶▶—UNLOCK Admin—adminname————▶▶
```

Parameters

adminname

Specifies the name of the administrator to be unlocked. This parameter is required.

Examples

Task

The administrator JOE is locked out of ADSM. Permit JOE to re-access the server.

Command: UNLOCK Admin joe

Output:

```
ANR2074I Administrator JOE unlocked.
```

Related Commands

Table 141 provides references to related commands.

Table 141. Commands Related to UNLOCK ADMIN

Command	Description	Reference
LOCK ADMIN	Prevents an administrator from accessing ADSM.	Page 244
QUERY ADMIN	Displays information about one or more ADSM administrators.	Page 256

UNLOCK NODE (Unlock a Client Node)

Use the UNLOCK NODE command to allow a locked client node to reaccess the server.

By using the LOCK NODE and UNLOCK NODE commands, you can temporarily restrict an administrator from accessing the server.

Privilege Class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the client node is assigned.

Syntax

```
▶▶ UNLOCK Node nodename ▶▶
```

Parameters

nodename

Specifies the name of the client node to be unlocked. This parameter is required.

Examples

Task

The client node LARRY is locked out of ADSM. Permit LARRY to reaccess the server.

Command: UNLOCK Node larry

Output:

```
ANR2066I Node LARRY unlocked.
```

Related Commands

Table 142 provides references to related commands.

Table 142. Commands Related to UNLOCK NODE

Command	Description	Reference
LOCK NODE	Prevents a client from accessing the server.	Page 245
QUERY NODE	Displays partial or complete information about one or more clients.	Page 320

UPDATE ADMIN (Update an Administrator)

Use the UPDATE ADMIN command to change the password or contact information for an administrator. However, you cannot update the server console administrator name.

You can use this command to modify selected parameters for the specified administrator. If you do not explicitly update a parameter, it remains unchanged.

You must use the RENAME ADMIN command to change the name of a registered administrator.

Passwords for both administrative and client node sessions must be changed after a length of time determined by the password expiration period. The SET PASSEXP command is used to set this period.

Privilege Class

To issue this command to change another administrator's password or contact information, you must have system privilege. Any administrator can issue this command to update his or her own password or contact information.

Syntax

```

▶▶—UPDate Admin—(1)—adminname—password—CONtact=contactinfo—▶▶

```

Note:

¹ You must specify at least one optional parameter on this command.

Parameters

adminname

Specifies the name of the administrator to be updated. This parameter is required.

password

Specifies the administrator's password. This parameter is optional. Passwords remain current for a period determined by the password expiration period.

CONtact=*contactinfo*

Specifies a text string that identifies the administrator. This parameter is optional. The text string must be enclosed in quotation marks if it contains any blanks. To remove previously defined contact information, specify a null string ("") for this value.

Examples

Task

Update the administrator LARRY to have the password SECRETWORD.

Command: UPDate Admin larry secretword

Output:

```
ANR2071I Administrator LARRY updated.
```

Related Commands

Table 143 provides references to related commands.

Table 143. Commands Related to UPDATE ADMIN

Command	Description	Reference
QUERY ADMIN	Displays information about one or more ADSM administrators.	Page 256
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.	Page 342
RENAME ADMIN	Changes an ADSM administrator's name.	Page 380
SET PASSEXP	Specifies the number of days after which a password is expired and must be changed.	Page 423

UPDATE COPYGROUP (Update a Copy Group)

Use the UPDATE COPYGROUP command to update a defined backup or archive copy group.

You can use this command to modify selected parameters for the specified copy group. If you do not explicitly update a parameter, it remains unchanged.

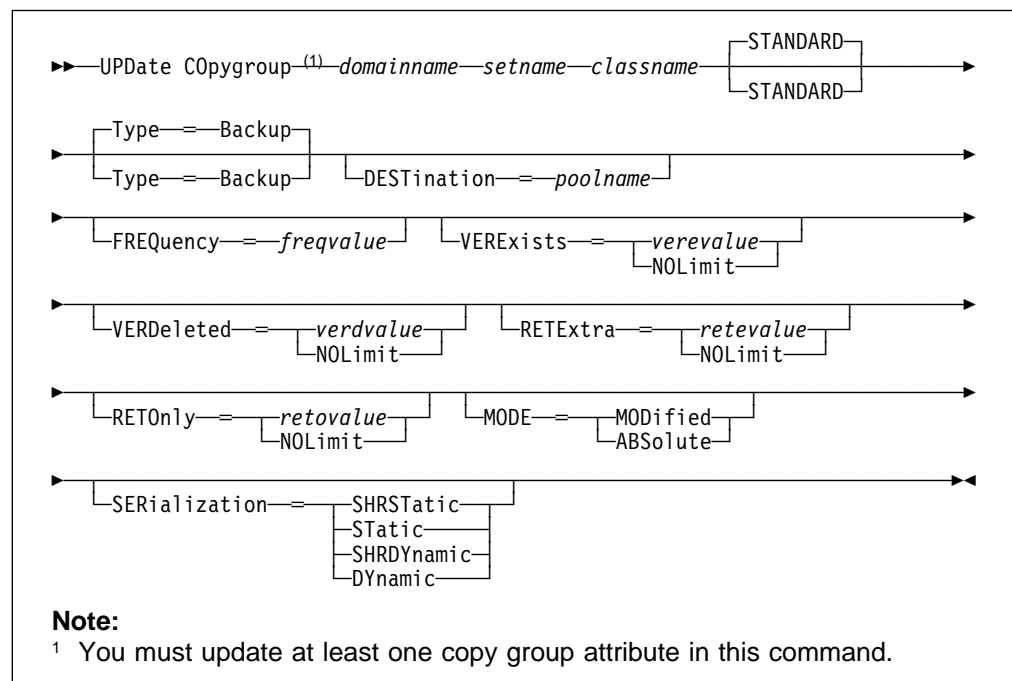
The UPDATE COPYGROUP command will fail if a copy storage pool name is specified as a destination.

The UPDATE COPYGROUP command takes two forms, depending upon whether the update is for a backup copy group or for an archive copy group. The syntax for each form is defined separately below.

Privilege Class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the copy group belongs.

Syntax for Updating a Backup Copy Group



Parameters for Updating a Backup Copy Group

domainname

Specifies the name of the policy domain in which the copy group is defined. This parameter is required.

setname

Specifies the name of the policy set in which the copy group to be updated is defined. This parameter is required.

UPDATE COPYGROUP

classname

Specifies the name of the management class to which the copy group belongs. This parameter is required.

STANDARD

Specifies the name of the copy group. The name of the copy group must be STANDARD. The default value is STANDARD.

Type=Backup

Specifies that you want to update a backup copy group. This parameter is optional. The default value is BACKUP.

DESTination=poolname

Specifies the name of the primary storage pool in which backup data is to be stored. This parameter is optional. You can only specify one storage pool as the destination for a copy group.

FREQuency=freqvalue

Specifies the minimum number of days that must elapse between successive backups. This parameter is optional. ADSM does not perform a backup unless the specified number of days has elapsed since the last backup. You can specify a value from 0–9999, inclusive.

Backup frequency works in conjunction with the MODE parameter, which is described later.

VERExists=vervalue

Specifies the maximum number of different backup versions to retain for files that currently exist on the client file system. This parameter is optional.

If the maximum number of versions allowed is 5, and backup processing creates a sixth version, ADSM deletes the oldest backup version from server storage. Possible values are:

vervalue

Specifies the number of backup versions to retain for files that currently exist on the client file system. You can specify a value from 1–9999, inclusive.

NOLimit

Specifies that you want ADSM to retain all backup versions.

VERDeleted=verdvalue

Specifies the maximum number of backup versions to retain for files that have been deleted from the client file system after being backed up to the server. This parameter is optional. ADSM ignores this parameter if the file still resides on the client node.

If the client file is deleted, during the next backup ADSM changes the active backup version to inactive and erases the oldest versions that are in excess of the number specified by this parameter. The expiration date for the remaining versions is determined by the values specified for the RETEXTRA and RETONLY parameters. Possible values are:

verdvalue

Specifies the number of backup versions to retain for files that have been deleted from the client file system after being backed up to the server. You can specify a value from 0–9999, inclusive.

NOLimit

Specifies that you want ADSM to retain all backup versions for files that are deleted from the client file system after being backed up to the server.

RETE*Extra=retevalue*

Specifies the retention time, in days, for all but the most recent backup version. The value of this parameter determines which versions are deleted during inventory expiration processing. This parameter is optional. Possible values are:

retevalue

Specifies an integer that defines the number of days to retain all backup versions other than the most recent version. You can specify a value from 0–9999, inclusive. The default value is 30 days.

NOLimit

Specifies that you want ADSM to retain all backup versions, other than the most recent version, indefinitely (that is, inventory expiration processing does not affect the backup versions). The most recent version of a file that exists on the client workstation is retained indefinitely or until the user deletes the file on the workstation. If the user deletes the file on the workstation, the file is rolled off using the value specified for the VERDELETED parameter.

If the NOLIMIT value is specified, extra backup versions are deleted during backup processing under control of the VEREXISTS or VERDELETED parameters.

RETO*Only=retovalue*

Specifies the retention time, in days, that ADSM retains the last copy it has of a file that has been deleted from the client file system. This parameter is optional. Possible values are:

retovalue

Specifies the number of days to retain the last remaining inactive copy of a file. You can specify a value from 0–9999, inclusive.

NOLimit

Specifies that the last remaining inactive copy of a file be kept indefinitely.

If you specify NOLIMIT, extra backup versions are deleted under control of the VEREXISTS or VERDELETED parameters.

MODE*=mode*

Specifies whether a file should be backed up based on changes made to the file since the last time it was backed up or regardless of whether it has changed. This parameter is only used by incremental backup. This parameter is optional. Possible values are:

MOD*ified*

Specifies that the file is backed up only if it has changed since the last backup. A file is considered changed if any of the following is true:

- The date last modified is different
- The file size is different
- The file owner is different

- The file permissions are different

ABSolute

Specifies that the file is backed up regardless of whether it has been changed.

SERialization=*serialvalue*

Specifies whether a file or directory can be in use during a backup operation and what the server should do if it is in use. This parameter is optional. The default value is SHRSTATIC.

Warning: Be careful about choosing a management class that contains a copy group that specifies SERIALIZATION=SHRDYNAMIC or SERIALIZATION=DYNAMIC.

If you choose a management class that allows a file to be backed up while it is in use, the backup version stored on the server might be a fuzzy backup. A fuzzy backup is a backup version that may not accurately reflect what is currently in the file. It might contain some, but not all, changes. If that is not acceptable, be sure to choose a management class that allows ADSM to create a backup version only if the file is not being used.

If a file is restored that contains a fuzzy backup, the file might not be usable, depending on the file's application.

Possible values are:

SHRStatic

Specifies that a file or directory must not be modified during backup. The server attempts to perform a backup as many as four times, depending on the value specified for CHANGINGRETRIES specified in the client user options file. If the file or directory is in use during the backup process, the server will not back it up.

Static

Specifies that a file or directory must not be modified during backup. If the file or directory changes during the backup attempt, it is not backed up.

Platforms that do not support the STATIC option default to SHRSTATIC. The STATIC option is not supported on the DOS platform.

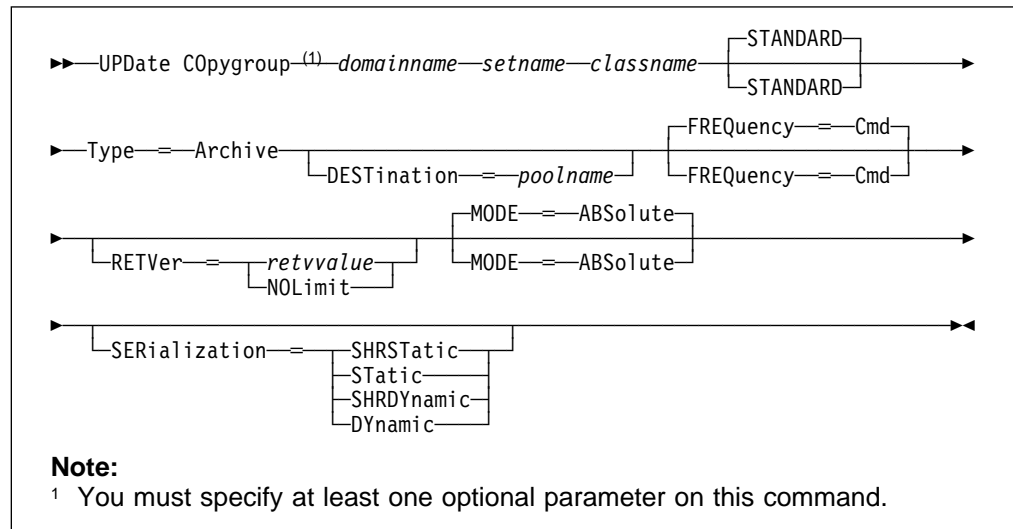
SHRDynamic

Specifies that if the file or directory changes during the backup attempt, the server will back it up on its last try even though the file or directory is in use.

Dynamic

Specifies that a file or directory is backed up on the first attempt by the server, regardless if the file or directory is in use during the backup process.

Syntax for Updating an Archive Copy Group



Parameters for Updating an Archive Copy Group

domainname

Specifies the name of the policy domain in which the copy group is defined. This parameter is required. The maximum length of this name is 30 characters.

setname

Specifies the name of the policy set in which the copy group belongs. This parameter is required. The maximum length of this name is 30 characters.

classname

Specifies the name of the management class to which the copy group belongs. This parameter is required. The maximum length of this name is 30 characters.

STANDARD

Specifies the name of the copy group. The name of the copy group must be STANDARD. The default value is STANDARD.

Type=Archive

Specifies that you want to update an archive copy group. This parameter is required.

DESTINATION=poolname

Specifies the name of the storage pool where the archive files are to reside. This parameter is optional. You can only specify on storage pool as the destination for a copy group.

FREQUENCY=Cmd

Specifies that a file archived. The copy frequency must be CMD. This parameter is optional. A file is archived only when specified in an ARCHIVE command or when selected for archive using the graphical user interface.

RETVer=retvvalue

Specifies the number of days to retain an archive copy that is bound to this management class. The value of this parameter determines whether a copy is deleted during inventory expiration processing. This parameter is optional.

retvvalue

Specifies the number of days to retain an archive copy. You can specify a value from 0–9999, inclusive.

NOLimit

Specifies that an archive copy is maintained indefinitely (that is, inventory expiration processing does not affect the archive copy).

MODE=ABSolute

Specifies that the file is archived regardless of whether it has been changed since the last archive request. This parameter is optional. The default value is ABSOLUTE.

SERialization=serialvalue

Specifies whether a file can be in use during an archive operation and what the server should do if it is in use. This parameter is optional. The default value is SHRSTATIC.

Warning: Be careful about choosing a management class that contains a copy group that specifies **SERIALIZATION=SHRDYNAMIC** or **SERIALIZATION=DYNAMIC**.

If you choose a management class that allows a file to be archived while it is in use, the archive copy stored on the server might be a fuzzy copy. A fuzzy copy is an archive copy that may not accurately reflect what is currently in the file. It might contain some, but not all, changes. If that is not acceptable, be sure to choose a management class that allows ADSM to create an archive copy only if the file is not being used.

If a file is restored that contains a fuzzy copy, the file might not be usable, depending on the file's application.

Possible values are:

SHRStatic

Specifies that a file must not be modified during archive. The server attempts to perform an archive as many as four times, depending on the value specified for **CHANGINGRETRIES** specified in the client user options file. If the file is in use during the archive process, the server will not archive.

Static

Specifies that a file must not be modified during archive. If the file changes during the archive attempt, it is not archived.

Platforms that do not support the **STATIC** option default to **SHRSTATIC**. The **STATIC** option is not supported on the DOS platform.

SHRDynamic

Specifies that if the file changes during the archive attempt, the server will archive it on its last try even though the file is in use.

Dynamic

Specifies that a file is archived on the first attempt by the server, regardless if the file is in use during the archive process.

Examples

Task 1

Update an archive copy group named STANDARD in the PROG1 policy domain, SUMMER policy set, MCLASS1 management class. The archive destination is ARCHPOOL, where the archive copies are kept for 190 days.

Command: UPDate C0pygroup prog1 summer mclass1 STANDARD Type=Archive
DESTINATION=archpool RETVer=190

Output:

```
ANR1537I Archive copy group STANDARD updated in policy domain
PROG1, set SUMMER, management class MCLASS1.
```

Task 2

Update a backup copy group named STANDARD in the EMPLOYEE_RECORDS policy domain, VACATION policy set, ACTIVEFILES management class. The backup destination is BACKPOOL, with a minimum interval of seven days between backups, regardless of whether the files have been modified. Retain up to two active backup versions.

Command: UPDate C0pygroup employee_records vacation activefiles STANDARD
Type=Backup DESTINATION=backpool FREQUENCY=7 VEREXISTS=2
MODE=ABSolute

Output:

```
ANR1532I Backup copy group STANDARD updated in policy domain
EMPLOYEE_RECORDS, set VACATION, management class ACTIVEFILES.
```

Related Commands

Table 144 provides references to related commands.

Table 144 (Page 1 of 2). Commands Related to UPDATE COPYGROUP

Command	Description	Reference
ASSIGN DEFMGMTCLASS	Assigns an existing management class as the default for a specified policy set.	Page 40
COPY MGMTCLASS	Creates a copy of an existing management class.	Page 73
DEFINE COPYGROUP	Defines a copy group for backup or archive processing within a specified management class.	Page 83

UPDATE COPYGROUP

Table 144 (Page 2 of 2). Commands Related to UPDATE COPYGROUP

Command	Description	Reference
DEFINE MGMTCLASS	Defines a management class within the policy domain and policy set specified.	Page 114
DELETE COPYGROUP	Deletes a backup or archive copy group from a policy domain and policy set.	Page 153
DELETE MGMTCLASS	Deletes a management class and its copy groups from a policy domain and policy set.	Page 172
EXPIRE INVENTORY	Manually starts inventory expiration processing.	Page 193
QUERY COPYGROUP	Displays the attributes of an existing copy group.	Page 270
QUERY MGMTCLASS	Displays information about one or more management classes.	Page 314

UPDATE DBBACKUPTRIGGER (Update the Database Backup Trigger)

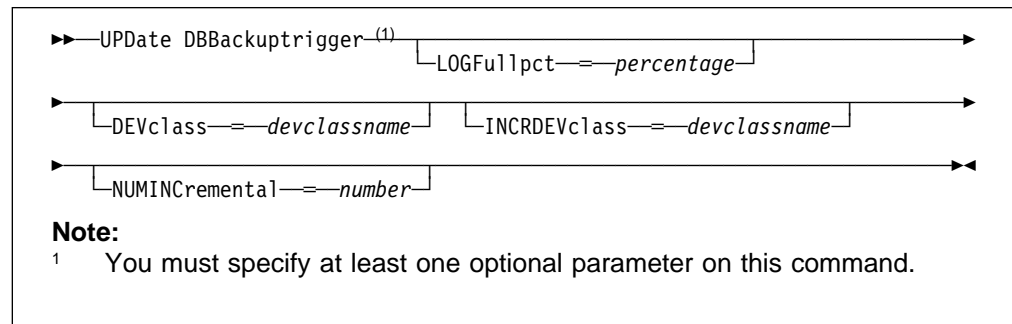
Use the UPDATE DBBACKUPTRIGGER command to change one or more parameters for the database backup trigger. You can use this command to change selected parameters. Any parameters you do not specify remain unchanged.

To display the current settings, use the QUERY DBBACKUPTRIGGER command.

Privilege Class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax



Parameters

LOGFullpct=*percentage*

Specifies when ADSM automatically backs up the database, based on the percentage of available space used for the recovery log. When the percentage of space used for the recovery log exceeds this value, ADSM performs an automatic backup of the database and removes unnecessary records from the recovery log. This parameter is optional. You can specify a value from 1–99. The default value is 50 percent.

DEVclass=*devclassname*

Specifies the name of the sequential access device class to use for backups. This parameter is optional.

If the mount limit for the device class specified has already been reached (all drives are busy) when the backup runs, ADSM automatically cancels lower priority operations, such as reclamation, to make a mount point available for the backup.

Be sure that you have used the DEVCONFIG option in the server options file to specify an external file in which to store a backup copy of device class definitions. Otherwise, if your ADSM database is damaged or lost and must be restored, the definitions created using the DEFINE DEVCLASS command will not be available and will need to be recreated manually.

You should also use the VOLUMEHISTORY option in the server options file, to specify an external file in which to store a backup copy of sequential volume history information. The volume history file is also used for database recovery.

UPDATE DBBACKUPTRIGGER

INCRDEVclass=*devclassname*

Specifies the name of the sequential access device class to use for incremental backups. This parameter is optional. If you do not specify a value for this option, all backups are run using the device class specified with the DEVCLASS parameter.

If the mount limit for the device class specified has already been reached (all drives are busy) when the backup runs, ADSM automatically cancels lower priority operations, such as reclamation, to make a mount point available for the backup.

Be sure that you have used the DEVCONFIG option in the server options file to specify an external file in which to store a backup copy of device class definitions. Otherwise, if your ADSM database is damaged or lost and must be restored, the definitions created using the DEFINE DEVCLASS command will not be available and will need to be recreated manually.

NUMINCremental=*number*

Specifies the maximum number of incremental backups that can be run before a full backup is required. This parameter is optional. You can specify a value from 0–32. A 0 specifies that ADSM runs only full backups automatically. The default value is 6.

Each incremental backup, whether run automatically by ADSM or run by using the BACKUP DB command, is added to the count of incremental backups run.

Each full backup, whether run automatically by ADSM or run by using the BACKUP DB command, resets the count for incremental backups to 0.

Examples

Task

Reset the database backup trigger at 60 percent and the number of incremental backups that ADSM runs between full backups to 5.

Command: `UPDate DBBackuptrigger LOGFull=60 NUMINCremental=5`

Output:

```
ANR2285I Database backup trigger updated, but disabled.
```

Related Commands

Table 145 provides references to related commands.

Table 145 (Page 1 of 2). Commands Related to UPDATE DBBACKUPTRIGGER

Command	Description	Reference
BACKUP DB	Backs up the ADSM database to sequential access volumes.	Page 50
DEFINE DBBACKUPTRIGGER	Defines a trigger for automatic ADSM database backups.	Page 91

Table 145 (Page 2 of 2). Commands Related to UPDATE DBBACKUPTRIGGER

Command	Description	Reference
DEFINE DEVCLASS	Defines a device class that can be used for storage pools, importing, or exporting.	Page 99
DELETE DBBACKUPTRIGGER	Deletes the database backup trigger.	Page 156
QUERY DBBACKUPTRIGGER	Displays information about the database backup trigger.	Page 278
SET LOGMODE	Sets the mode for saving records in the recovery log.	Page 416

UPDATE DEVCLASS (Change the Attributes of a Device Class)

Use the UPDATE DEVCLASS command to update a defined device class.

You can use this command to modify selected parameters for the specified device class. If you do not explicitly update a parameter, it remains unchanged.

At installation, ADSM provides two predefined device classes with the following device names and types: DISK and CARTRIDGE. The DISK device class cannot be modified or deleted.

Use this command to modify the CARTRIDGE device class, as well as any device classes that you have defined using the DEFINE DEVCLASS command.

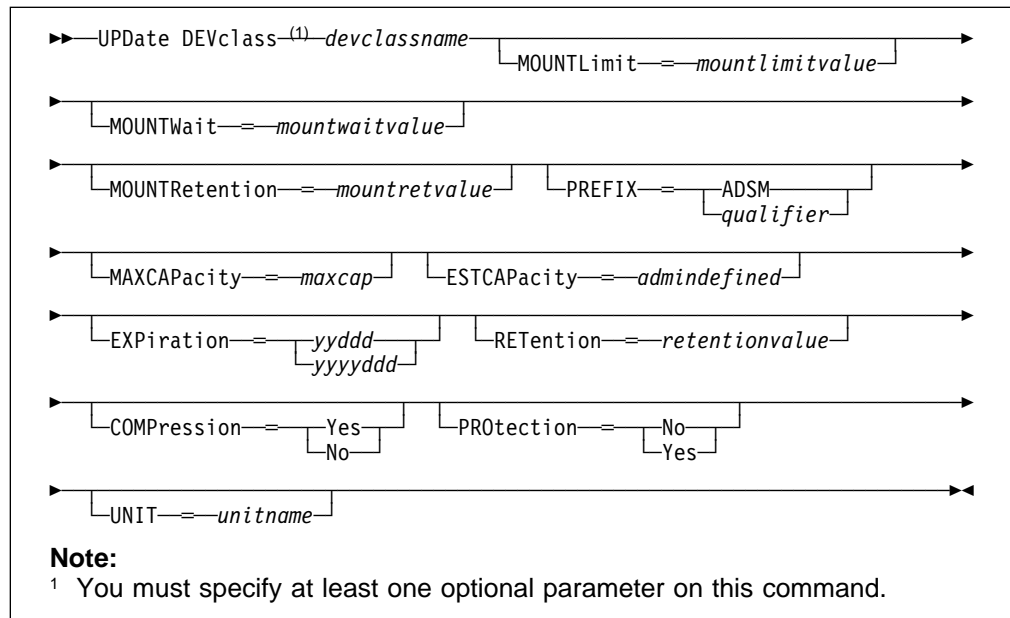
If you include the DEVCONFIG option in the server options file, the files you specify with that option are automatically updated with the results of this command. When you use this option, the files specified are automatically updated whenever a device class is defined, updated, or deleted.

The supported parameters for this command vary across device type and the environment in which your system runs. The syntax diagrams are shown using the complete syntax for the command according to the environment in which your system runs.

Privilege Class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax for Updating a Device Class with a Device Type of Cartridge



Parameters for Updating a Device Class With a Device Type of CARTRIDGE

devclassname

Specifies the name of the device class to be defined. This parameter is required. The maximum length of the device class name is 30 characters.

MOUNTLimit=*mountlimitvalue*

Specifies the maximum number of tape volumes that can be simultaneously mounted for a device class. This parameter is optional. You can specify an integer in the range of 1 to 256, inclusive.

If this value is set to 1, then it may not be possible for tape-to-tape operations (like reclamation) to work.

You cannot share drives between device classes (that is, the sum of your mount limits should not be more than the number of physical drives allocated to ADSM).

MOUNTWait=*mountwaitvalue*

Specifies the maximum length of time, in minutes, to wait for a tape volume to be mounted. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. You can specify an integer in the range of 0 to 9999, inclusive.

MOUNTRetention=*mountretvalue*

Specifies the amount of time, in minutes, to retain an idle tape volume before dismounting it. Mount retention time begins after the idle time-out period has expired. This parameter is optional. You can specify an integer in the range of 0 to 9999, inclusive.

This parameter can improve response time for tape mounts by leaving the previously mounted volume online.

PREFIX=*tapevolprefix*

Specifies the high-level qualifier of the data set name in the tape labels. For all tape volumes assigned to this device class, ADSM uses this prefix to create the data set name. This parameter is optional.

An example of a tape volume data set name using the default prefix is ADSM.BFS. You can specify any valid MVS data set qualifier for the prefix.

If you have already established a tape label naming convention that supports your current tape management system, consider using a tape volume prefix that conforms to your naming conventions.

MAXCAPacity=*maxcap*

Specifies the maximum amount of data ADSM can store in the volumes categorized by this device class. This parameter is optional.

You must specify this value as an integer followed by a scale factor. The following are possible scale factors:

Symbol	Symbol Name
K	Kilobyte
M	Megabyte
G	Gigabyte

For example, MAXCAPACITY=5M specifies that the maximum capacity for this device class is 5 megabytes.

ESTCAPacity=*admindefined*

Specifies the estimated capacity for volumes categorized by this device class. This parameter is optional.

This value is ignored if you specify a value other than 0 in the MAXCAPACITY parameter. If ADSM uses the value specified in the ESTCAPACITY parameter, this value does not determine the amount of data stored on the volume, but only determines the usage before the volume is filled. After a volume is full, the actual amount of data stored on the tape is used for these calculations.

You must specify this value as an integer followed by a scale factor. Possible scale factors are:

Symbol	Symbol Name
K	Kilobyte
M	Megabyte
G	Gigabyte

For example, ESTCAPACITY=5M specifies that the estimated capacity for a volume in this device class is 5 megabytes. The minimum value accepted is 100KB (that is, ESTCAPACITY=100K).

For more information on the default estimated capacity for cartridge tapes, see Table 146.

Table 146. ADSM Estimated Capacity for Cartridge Tape

Tape Format	Media Type	Estimated Capacity
3480	(Not applicable)	180MB
3480XF	(Not applicable)	180MB
3490-A01	CST	180MB
3490B	CST	360MB
3490C	CST	360MB
3490B	ECCST	720MB
3490C	ECCST	720MB

Note:

3490 Base devices include models A01, A02, A10, A20, B02, B04, B20, B40, D41, D42
 3490 E devices include models C1A, C2A, C10, C11, C22

EXPIration=*expirationdate*

Specifies the expiration date that is placed on the tape labels for this device class, which are used for allocation. This parameter is optional.

If you specify this parameter, enter the expiration date when ADSM no longer requires the tape. ADSM does not use this information, however; this information is passed to the operating system for use by MVS or tape management systems. For additional information, see the *ADSM Administrator's Guide*.

If you specify this parameter, you cannot specify the RETENTION parameter. Possible values are:

yyddd

Specifies the expiration date for this device class, in YYDDD format.

yyyyddd

Specifies the expiration date for this device class, in YYYYDDD format.

RETention=*retentionvalue*

Specifies an integer that defines the length of time, in days, to retain the tape. This parameter is optional. You can specify an integer in the range of 1 to 9999, inclusive. If you specify this parameter, enter the number of days that equals the length of time ADSM is expected to use the tape. ADSM does not use this information, however; this information is passed to the operating system for use by MVS or tape management systems. For additional information, see the *ADSM Administrator's Guide*.

If you specify this parameter, you cannot specify the EXPIRATION parameter.

COMPression=*compressionvalue*

Specifies whether the device class is to exploit file compression. This parameter is optional.

If your tape devices have IDRC, then data can be compressed when written to CARTRIDGE tape devices. If your tape devices do not have IDRC, then the compression attribute is ignored and data is not compressed.

Also, if your tape volumes can be moved to other physical sites, ensure that the compressed data can be read at the other sites. If the other sites cannot read compressed data, you want to specify COMPRESSION=NO. Possible values are:

Yes

Specifies that the data for each tape volume is compressed.

No

Specifies that the data for each tape volume is not compressed.

PROtection=*protectionvalue*

Specifies whether RACF, if installed, protects all volumes assigned to this device class. This parameter is optional. Possible values are:

No

Specifies that RACF does not protect all CARTRIDGE volumes assigned to this tape device.


Yes

Specifies that RACF protects all CARTRIDGE volumes assigned to this tape device.

UNIT=*unitname*

Specifies an esoteric unit name to specify a group of tape devices that support Cartridge System tape. This parameter is optional. This name can be as many as 8 characters.

Syntax for Updating a Device Class with a Device Type of FILE



Parameters for Updating a Device Class With a Device Type of FILE

devclassname

Specifies the name of the device class to be defined. This parameter is required. The maximum length of the device class name is 30 characters.

MOUNTLimit=*mountlimitvalue*

Specifies the maximum number of FILE volumes that can be concurrently open at any one time for this device class. This parameter is optional. For 3995 devices emulating 3390 devices, this value should not be set higher than the number of concurrent input/output streams possible on the media storing the ADSM volumes.

Because FILE volumes are classified as sequential media, processing conventions normally associated with sequential media are followed. Specifically, this means that mount limits are observed, and that data that has expired leaves *holes* in the volume and volumes can only be recycled through either the MOVE DATA command or the volume reclamation process. Only one process can read or write to a volume at a time. This has important implications for clients backing up directly to FILE media. The number of active clients is limited to the value specified in the MOUNTLIMIT parameter for the device class.

The value you specify in this parameter is important if there is a significant penalty switching from one volume to another (switching can take place when using 3995 devices emulating 3390 devices). The value you specify in this parameter should be no higher than the number of physical drives available on the device.

See the Usage Notes for more information on FILE volumes.

PREFIX=*qualifier*

Specifies the MVS high-level qualifier that is used to allocate scratch volume data sets. It must be 1–8 alphanumeric characters and begin with an alpha.

MAXCapacity=*maxcap*

Specifies the maximum capacity of the FILE volumes. This parameter is optional. The largest accepted value is 2047MB (2GB minus 1MB). The smallest accepted value is 100KB.

Volume capacity is determined by the value you specify for the device class. If you are explicitly allocating FILE volume data sets, be sure the allocated capacity matches the MAXCAPACITY parameter. FILE volumes are allocated in 16KB blocks. Thus, if the MAXCAPACITY value is 50MB, you should allocate 3200 blocks (16KB per block). A combination of primary and secondary extents may be specified to achieve the allocation.

Note: Note that FILE volumes will only be filled to the amount specified by the MAXCAPACITY parameter. Thus, if a larger data set is explicitly allocated, the additional space will not be used.

If the data set is allocated to a size smaller than the value you specify for the MAXCAPACITY parameter, writing will stop after all allocated blocks have been used. If this situation occurs, a message is issued, the volume status remains “Filling,” and access is set to read-only.

You must specify this value as an integer followed by a scale factor. Possible scale factors are:

Symbol	Symbol Name
K	Kilobyte
M	Megabyte
G	Gigabyte

For example, MAXCAPACITY=5M specifies that the maximum capacity for this device class is 5 megabytes.

Usage Notes

- The 3490-A01 is treated like a 3480. The capacity shown in a QUERY VOLUME will, therefore, reflect 3480 capacities.
- The estimated capacity of a tape volume is determined by the format, length, and density of the volume.
- The format for a tape volume is determined by the device type and the recording technology of the tape volume.
- The EXPIRATION and RETENTION parameters are mutually exclusive; if you specify one, you cannot specify the other.
- FILE device class volumes are represented by sequential MVS data sets. There are two primary methods for allocating FILE volumes: scratch and explicit allocation.

ADSM allocates scratch volumes when you specify the MAXSCRATCH parameter on a storage pool allocated from a FILE device class. ADSM generates the name of your data set using a combination of the value you specified in the PREFIX parameter as the HLQ (high-level qualifier) and a system-generated extension.

Judicious use of the PREFIX parameter can assist the installation in identifying and isolating ADSM FILE volumes.

You can also use the DEFINE VOLUME command to explicitly define volumes to storage pools from a FILE device class. The volume name you specify is used as the fully qualified data set representing the volume. Explicitly defined volumes are used before scratch volumes. The data set representing the explicitly defined volume may or may not be allocated ahead of time. If it is not allocated in advance, the first time write processing occurs ADSM attempts to dynamically allocate the file.

The QUERY VOLUME command lists the fully qualified data set name corresponding to the FILE volume.

- ADSM attempts to dynamically allocate a data set for any volume not explicitly allocated ahead of time (this would be the usual way of allocating scratch volumes). If the capacity of a volume is 256MB or less, a single primary extent (0 secondary extents) is allocated. For volumes greater than 256MB, a primary extent of 256MB is allocated, with secondary extents of 1/15th of the remaining required space.

Explicitly defined volumes may either be defined before they are used or dynamically allocated. If the volume is predefined, it should have the following characteristics:

UPDATE DEVCLASS

```
Data Set Name: ADSM.FILEVOL0

      General Data:
Volume serial:      volser
Device type:        3390
Organization:       PS
Record format:      F
Record length:      16384
Block size:         16384
* First extent blocks: 16386
* Secondary blocks: 1093

Creation date:      1995/02/16
Expiration date:    ***NONE***
```

* These values vary with volume size.

- FILE volume reclamation uses the same methodology as other sequential media. The key factor for recycling a reclaimed volume is whether the volume was automatically defined by ADSM as a scratch volume or explicitly defined by an ADSM administrator using the DEFINE VOLUME command. Explicitly defined volumes are marked as empty and available for subsequent mounts at the completion of reclamation. Scratch volumes are deleted from the storage pool in which it was defined and the MVS data set corresponding to the scratch volume is automatically erased from the system.

Examples

Task 1

Update the amount of time to retain idle mounted volumes for the ROUND1 device class to 75 minutes.

Command: UPDATE DEVCLASS ROUND1 MOUNTRETENTION=75

Output:

```
ANR2212I Device class ROUND1 updated.
```

Task 2

Update the SQUARE1 device class to have RACF protection and a maximum mount limit of 1.

Command: UPDATE DEVCLASS SQUARE1 MOUNTLIMIT=1 PROTECTION=Yes

Output:

```
ANR2205I Device class SQUARE1 updated.
```


Related Commands

Table 147 provides references to related commands.

Table 147. Commands Related to UPDATE DEVCLASS

Command	Description	Reference
BACKUP DEVCONFIG	Backs up ADSM device information to a file.	Page 54
DEFINE DEVCLASS	Defines a device class that can be used for storage pools, importing, or exporting.	Page 99
DELETE DEVCLASS	Deletes a device class name.	Page 160
QUERY DEVCLASS	Displays information about one or more device classes.	Page 283

UPDATE DOMAIN (Update a Policy Domain)


Use the UPDATE DOMAIN command to modify an existing policy domain.

You can use this command to update selected parameters for the specified policy domain. Any parameters that are not explicitly updated remain unchanged.

Privilege Class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the specified policy domain.

Syntax



Parameters

domainname

Specifies the name of an existing policy domain to be updated. This parameter is required. The maximum length of this name is 30 characters.

DESCRIPTION=*description*

Describes the policy domain using a text string. This parameter is optional. The maximum length of the description is 255 characters. The description must be enclosed in quotation marks if it contains any blank characters. To remove a previously defined description, specify a null string ("") for this value.

BACKRETENTION=*bkretvalue*

Specifies the number of days (from the date of deactivation) to retain inactive backup versions. The backup retention value is used by ADSM to manage inactive copies of files when either of the following conditions are present:

- A file's management class is changed, but neither the new management class or default management class contains a backup copy group.
- The management class to which a file is bound no longer exists, and the default management class does not contain a backup copy group.

ARCHRETENTION=*archretvalue*

Specifies the number of days (from the date of archive) to retain archive copies. The archive retention value is used by ADSM to manage inactive copies of files when either of the following conditions are present:

- A file's management class is changed, but neither the new management class or default management class contains an archive copy group.
- The management class to which a file is bound no longer exists, and the default management class does not contain an archive copy group.

Examples

Task 1

Update the description of the EMPLOYEE_RECORDS policy domain to Employee Information and set the backup retention parameter to 50 days.

Command: UPDate D0main employee_records DESCription="Employee Information" BACKRETention=50

Output:

```
ANR1502I Policy domain EMPLOYEE_RECORDS updated.
```

Task 2

Update the description of the PROG1 policy domain to Development Programming Group. Set the backup retention parameter to 45 days and the archive retention parameter to 75 days.

Command: UPDate D0main prog1 DESCription="Development Programming Group" BACKRETention=45 ARCHRETention=75

Output:

```
ANR1502I Policy domain PROG1 updated.
```

Related Commands

Table 148 provides references to related commands.

Table 148. Commands Related to UPDATE DOMAIN

Command	Description	Reference
COPY DOMAIN	Creates a copy of an existing policy domain.	Page 70
DEFINE DOMAIN	Defines a policy domain that can be assigned to one or more clients.	Page 107
DEFINE POLICYSET	Defines a policy set within the specified policy domain.	Page 118
DELETE DOMAIN	Deletes a policy domain along with any policy objects in the policy domain.	Page 162
QUERY DOMAIN	Displays information about one or more policy domains.	Page 287

UPDATE MGMTCLASS (Update a Management Class)

Use the UPDATE MGMTCLASS command to modify an existing management class. You can use this command to update selected parameters for the specified management class. Any parameters not explicitly updated remain unchanged.

The UPDATE MGMTCLASS command fails if a copy storage pool is specified as the destination for space-managed files.

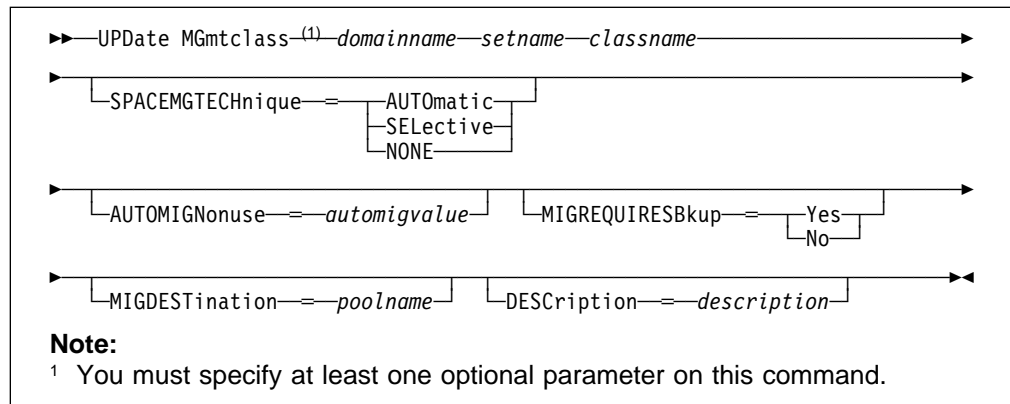
The following parameters affect only HSM clients that use this management class, and have no effect on backup-archive clients:

SPACEMGTECHNIQUE
 AUTOMIGNONUSE
 MIGREQUIRESBKUP
 MIGDESTINATION

Privilege Class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the policy set belongs.

Syntax



Parameters

domainname

Specifies the name of the policy domain to which the management class belongs. This parameter is required. The maximum length of this name is 30 characters.

setname

Specifies the name of the policy set to which the management class belongs. This parameter is required. The maximum length of this name is 30 characters. You cannot update a management class that belongs to the ACTIVE policy set.

classname

Specifies the name of the management class that you want to update. This parameter is required.

SPACEMGTECH*nique=smtvalue*

Specifies the space management technique to use for the management class. This parameter is optional. Possible values are:

AUT*OMATIC*

Specifies that the file is eligible for automatic migration or selective migration.

SE*LECTIVE*

Specifies that the file is eligible for selective migration only.

NONE

Specifies that the file is not eligible for migration.

The default in the management class named STANDARD is NONE.

This parameter is effective only for HSM clients, not for backup-archive clients.

AUTOMIG*Nonuse=automigvalue*

Specifies the number of days that must elapse since the file was last used before it is eligible for automatic migration. This parameter is optional.

If the SPACEMGTECHNIQUE parameter is not AUTOMATIC, ADSM ignores this attribute. You can specify a value from 0–9999. The default value is 0.

This parameter is effective only for HSM clients, not for backup-archive clients.

MIGREQUIRES*Bkup=migreqvalue*

Specifies whether a backup version of a file must exist before a file can be migrated. This parameter is optional. Possible values are:

Yes

A backup version must exist.

No

A backup version is optional.

This parameter is effective only for HSM clients, not for backup-archive clients.

MIGDEST*ination=poolname*

Specifies the name of a primary storage pool that is a destination for migrated files. This parameter is optional. The default destination is SPACEMGPOOL.

A file must be associated with a management class that specifies an appropriate storage pool for migrated files. The type of storage devices in a storage pool and options specified for the storage pool affects the speed with which a migrated file can be recalled to a local file system.

The command fails if a copy storage pool is specified as the destination.

This parameter is effective only for HSM clients, not for backup-archive clients.

DESC*ription=description*

Specifies a text string that describes the management class. This parameter is optional. The maximum length of the description is 255 characters. The description must be enclosed in quotation marks if it contains any blank char-

UPDATE MGMTCLASS

acters. To remove a previously defined description, specify a null string ("") for this value.

Examples

Task

Update the management class called MCLASS1, for policy set SUMMER in the PROG1 policy domain, with a description of Technical Design Support Mgmt Class.

Command: UPDate MGmtclass prog1 summer mclass1 DEScription="Technical Design Support Mgmt Class"

Output:

```
ANR1522I Management class MCLASS1 updated in policy domain
PROG1, set SUMMER.
```

Related Commands

Table 149 provides references to related commands.

Table 149 (Page 1 of 2). Commands Related to UPDATE MGMTCLASS

Command	Description	Reference
ASSIGN DEFMGMTCLASS	Assigns an existing management class as the default for a specified policy set.	Page 40
COPY MGMTCLASS	Creates a copy of an existing management class.	Page 73
DEFINE COPYGROUP	Defines a copy group for backup or archive processing within a specified management class.	Page 83
DEFINE POLICYSET	Defines a policy set within the specified policy domain.	Page 118
DEFINE MGMTCLASS	Defines a management class within the policy domain and policy set specified.	Page 114
DELETE MGMTCLASS	Deletes a management class and its copy groups from a policy domain and policy set.	Page 172
QUERY COPYGROUP	Displays the attributes of an existing copy group.	Page 270
QUERY MGMTCLASS	Displays information about one or more management classes.	Page 314

Table 149 (Page 2 of 2). Commands Related to UPDATE MGMTCLASS

Command	Description	Reference
QUERY POLICYSET	Displays information about one or more policy sets.	Page 329
UPDATE COPYGROUP	Changes one or more attributes of a copy group.	Page 439

UPDATE NODE (Update a Client Node)

Use the UPDATE NODE command to modify the attributes of a registered client node.

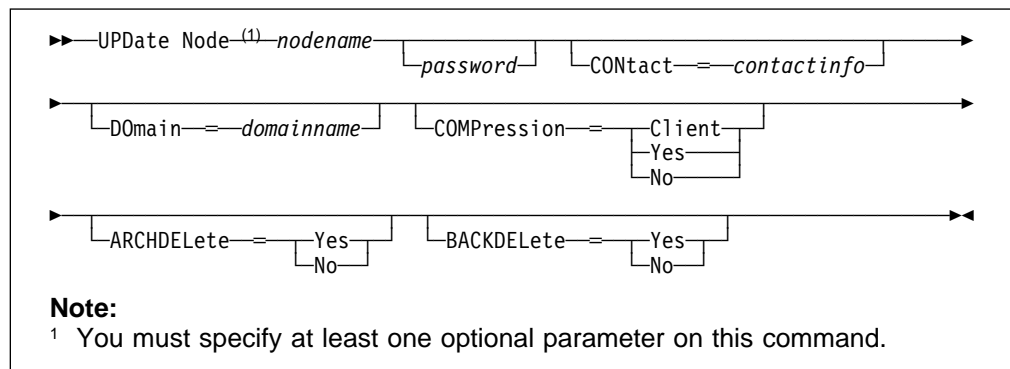
You can use this command to modify selected parameters for the specified client node. If you do not explicitly update a parameter, it remains unchanged.

You must use the RENAME NODE command to change the name of a registered client node.

Privilege Class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the client node belongs.

Syntax



Parameters

nodename

Specifies the name of the client node to be updated. This parameter is required.

password

Specifies the new password for the client node. This parameter is optional. The maximum length of the name is 64 characters. Passwords remain current for a period determined by the password expiration period.

CONTACT=contactinfo

Specifies a text string of information identifying the client node. This parameter is optional. The maximum length of the text string is 255 characters. The contact information must be enclosed in quotation marks if it contains any blanks. To remove previously defined contact information, specify a null string ("") for this value.

DOMAIN=domainname

Specifies the name of the policy domain to which the client node is registered. This parameter is optional.

COMPression=*compressionvalue*

Specifies whether the client node compresses its files before sending them to the server for backup and archive. This parameter is optional. Possible values are:

Client

Specifies that the client determines whether files will be compressed.

Yes

Specifies that the client node compresses its files before sending them to the server for backup and archive.

No

Specifies that the client node does not compress its files before sending them to the server for backup and archive.

ARCHDElete=*archdeletevalue*

Specifies whether the client node can delete its own archived files from the server. This parameter is optional. Possible values are:

Yes

Specifies that the client node can delete its own archive files from the server.

No

Specifies that the client node cannot delete its own archive files from the server.

BACKDElete=*backdeletevalue*

Specifies whether the client node can delete its own backup files from the server. This parameter is optional. Possible values are:

Yes

Specifies that the client node can delete its own backup files from the server.

No

Specifies that the client node cannot delete its own backup files from the server.

Examples

Task

Update node JOEOS2 so that the data on node JOEOS2 will be compressed when it is backed up or archived by ADSM and so that the client cannot delete archived files.

Command: UPDate Node joeos2 COMPression=Yes ARCHDElete=No

Output:

```
ANR2060I Node JOEOS2 registered in policy domain AIXTEST1.
```

Related Commands

Table 150 provides references to related commands.

Table 150. Commands Related to UPDATE NODE

Command	Description	Reference
QUERY NODE	Displays partial or complete information about one or more clients.	Page 320
REGISTER NODE	Defines a client to the server and sets options for that user.	Page 374
REMOVE NODE	Removes a client from the list of registered nodes for a specific policy domain.	Page 378

UPDATE POLICYSET (Update a Policy Set)

Use the UPDATE POLICYSET command to modify the description of an existing policy set.

You can use this command to modify the description for the specified policy set. If you do not explicitly update the DESCRIPTION parameter, it remains unchanged.

You cannot modify the description of the ACTIVE policy set.

Privilege Class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the policy set belongs.

Syntax

```
►► Update Policyset domainname setname DESCRIPTION=description ►►
```

Parameters

domainname

Specifies the name of the policy domain to which the policy set belongs. This parameter is required. The maximum length of this name is 30 characters.

setname

Specifies the name of the policy set to be updated. This parameter is required. The maximum length of this name is 30 characters. You cannot modify the attributes of the ACTIVE policy set.

DESCRIPTION=*description*

Specifies a text string that describes the policy set. This parameter is required. The maximum length of the description is 255 characters. The description must be enclosed in quotation marks if it contains any blank characters. To remove a previously defined description, specify a null string (") for this value.

Examples

Task

Update a policy set called VACATION for the EMPLOYEE_RECORDS policy domain with a description of Schedule Planning Information.

Command: UPDate Policyset employee_records vacation
DESCRiption="schedule planning information"

Output:

```
ANR1512I Policy set VACATION updated in policy domain  
EMPLOYEE_RECORDS.
```

Related Commands

Table 151 provides references to related commands.

Table 151. Commands Related to UPDATE POLICYSET

Command	Description	Reference
ACTIVATE POLICYSET	Validates and activates a policy set within a policy domain.	Page 38
COPY MGMTCLASS	Creates a copy of an existing management class.	Page 73
DEFINE DOMAIN	Defines a policy domain that can be assigned to one or more clients.	Page 107
DEFINE MGMTCLASS	Defines a management class within the policy domain and policy set specified.	Page 114
DEFINE POLICYSET	Defines a policy set within the specified policy domain.	Page 118
DELETE POLICYSET	Deletes a policy set, including its management classes and copy groups, from a policy domain.	Page 174
QUERY POLICYSET	Displays information about one or more policy sets.	Page 329
VALIDATE POLICYSET	Verifies and reports on conditions the administrator must consider before activating the policy set.	Page 501

UPDATE SCHEDULE (Update a Client or an Administrative Command Schedule)

Use the UPDATE SCHEDULE command to modify an existing schedule. The UPDATE SCHEDULE command takes two forms, depending on whether the schedule applies to client operations or administrative commands. The syntax and parameters for each form are defined separately.

Updating Client Schedules

Use the UPDATE SCHEDULE command to modify selected parameters for the specified schedule. If you do not explicitly update a parameter, it remains unchanged.

This command does not change the client associations that have been made to this schedule. Any clients that are associated with the original schedule execute the modified schedule.

A schedule determines:

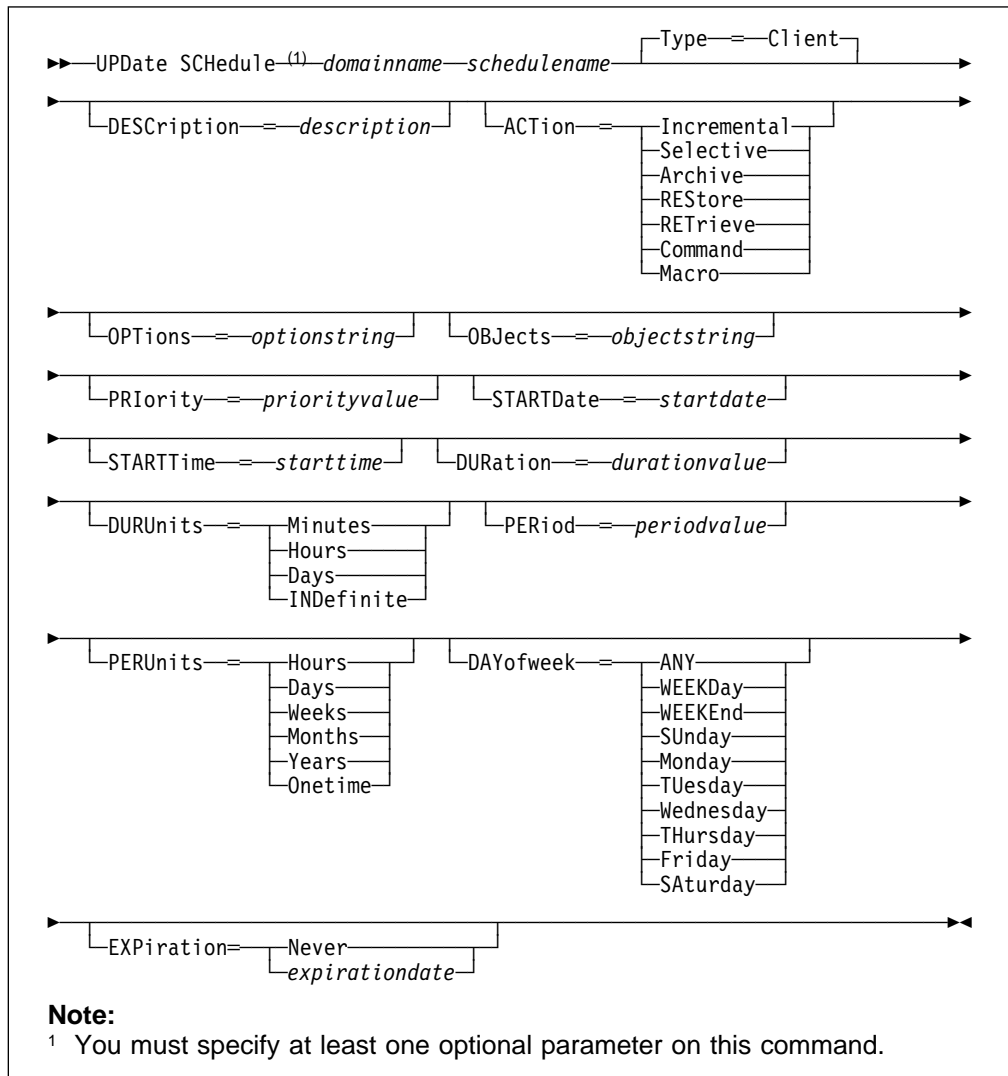
- The automatic function to be performed, for example, an incremental backup
- A startup window that indicates a time range during which the operation can begin
- The time period between startup windows

For each schedule, a startup window is specified. The startup window is the time period within which the scheduled operation must begin execution. The operation does not necessarily complete execution within this window. If the operation needs to be retried for any reason, the retry attempt must begin before the startup window elapses, or the operation is not retried.

Privilege Class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the schedule belongs.

Syntax for Updating Client Schedules



Parameters for Updating Client Schedules

domainname

Specifies the name of the policy domain to which this schedule belongs. This parameter is required.

schedulename

Specifies the name of the schedule to be updated. This parameter is required.

Type=Client

Specifies that this schedule updates a client schedule. This parameter is optional. The default value is CLIENT. Administrative command schedules are updated by specifying TYPE=ADMINISTRATIVE.

DESCRIPTION=description

Specifies a description of the schedule. This parameter is optional. The maximum length of the description is 255 characters. The description must be enclosed in quotation marks if it contains any blank characters. To

remove a previously defined description, specify a null string ("") for this value.

ACTION=*actionvalue*

Specifies the action that occurs when this schedule is executed. This parameter is optional. Possible values are:

Incremental

Specifies that you want this schedule to initiate the backup of any files that have changed since the last backup and for which the specified time between backups has elapsed.

Selective

Specifies that you want this schedule to back up only the files specified with the OBJECTS parameter.

Archive

Specifies that you want to archive the files specified with the OBJECTS parameter.

REStore

Specifies that you want this schedule to restore the files specified with the OBJECTS parameter.

REtrieve

Specifies that you want this schedule to retrieve the files specified with the OBJECTS parameter.

Command

Specifies that the action is to be a client operating system command string specified with the OBJECTS parameter.

Macro

Specifies that a macro whose file name is specified with the OBJECTS parameter is to be executed.

OPTions=*optionstring*

Specifies options that are supplied to the DSMC command at the time the schedule is executed. This parameter is optional. You can specify most options from the client's option file. For more information, refer to the appropriate client manual. Each option is preceded by a hyphen and delimited with a space. The option string must be enclosed in quotation marks if it contains any blank characters. When applicable, these options override the options specified by a client after it has successfully contacted the server.

Do not include the following options because they have no effect on the client's execution of the scheduled command:

MAXCMDRETRIES
 QUERYSCHEDPERIOD
 RETRYPERIOD
 SCHEDLOGNAME
 SCHEDMODE
 TCPCLIENTADDRESS
 TCPCLIENTPORT

OBJects=*objectstring*

Specifies the objects for which the specified action is performed. This parameter is optional. The allowed objects depend on the action you specify with

UPDATE SCHEDULE

the ACTION parameter and the operating system used by the clients associated with the schedule.

You can specify more than one object by separating the objects with commas, with no intervening spaces. For example, you can enter a string of file names or file spaces. If the value for this parameter contains any blank characters, you must enclose it in quotation marks.

When you specify ACTION=INCREMENTAL and do not specify any objects with this parameter, the DSMC command is invoked without specified objects and attempts to process objects as defined in the client's option file.

PRIority=*priorityvalue*

Specifies the priority value for a schedule. This parameter is optional. You can specify an integer from 1–10, inclusive, with 1 being the highest priority and 10 being the lowest.

If two or more schedules for a given client have the same window start time, this value determines the order of execution. The schedule with the highest priority starts first. For example, a schedule with PRIORITY=3 starts before a schedule with PRIORITY=5.

STARTDate=*startdate*

Specifies the date for the beginning of the window in which the schedule is first executed. This parameter is optional. This parameter is used in conjunction with the STARTTIME parameter to specify when the initial startup window begins.

Possible values are:

MM/DD/YYYY

Specifies the date for the beginning of the startup window in MM/DD/YYYY format (for example, 01/31/1995).

TODAY

Specifies that the beginning of the startup window is on the current date.

TODAY+*numdays* or +*numdays*

Specifies that the beginning of the startup window is a specific number of days after the current date. The maximum value for the number of days is 9999.

For example, to specify that the beginning of the initial startup window is 3 days from now, you can specify STARTDATE=TODAY+3 or STARTDATE=+3.

STARTTime=*starttime*

Specifies the time for the beginning of the window in which the schedule is first executed, in HH:MM format. This parameter is optional. This parameter is used in conjunction with the STARTDATE parameter to specify when the startup window begins.

DURation=*durationvalue*

Specifies the number of units that define the length of the startup window. This parameter is optional. You can specify an integer from 1–999, inclusive.

This parameter is used in conjunction with the DURUNITS parameter to specify the length of the startup window. For example, if DURATION=20 and DURUNITS=MINUTES, the schedule must be started within 20 minutes of the

start date and start time. The duration of the window must be shorter than the period between windows.

This value is ignored if you specify `DURUNITS=INDEFINITE`.

DURUnits=*durunitsvalue*

Specifies the time units used to determine the duration of the window in which this schedule can start. This parameter is optional.

This parameter is used in conjunction with the `DURATION` parameter to specify how long the startup window remains open to execute the schedule. For example, if `DURATION=20` and `DURUNITS=MINUTES`, the schedule must be started within 20 minutes of the start date and start time. The duration of the window must be shorter than the period between windows. Possible values are:

Minutes

Specifies that the duration of the window is defined in minutes.

Hours

Specifies that the duration of the window is defined in hours.

Days

Specifies that the duration of the window is defined in days.

INDefinite

Specifies that the startup window has an indefinite duration. The schedule can run any time after the scheduled start time, until the schedule expires. You cannot specify `DURUNITS=INDEFINITE` unless `PERUNITS=ONETIME` is also specified for the schedule.

PERiod=*periodvalue*

Specifies the length of time between startup windows for this schedule. This parameter is optional. You can specify an integer from 1–999, inclusive.

This parameter is used in conjunction with the `PERUNITS` parameter to specify the period between startup windows. For example, if you specify `PERIOD=5` and `PERUNITS=DAYS` (assuming that `DAYOFWEEK=ANY`), the operation is scheduled every 5 days after the initial start date and start time. The period between startup windows must exceed the duration of each window.

This value is ignored if you specify `PERUNITS=ONETIME`.

PERUnits=*perunitsvalue*

Specifies the time units used to determine the period between startup windows for this schedule. This parameter is optional.

This parameter is used in conjunction with the `PERIOD` parameter to specify the period between startup windows. For example, if `PERIOD=5` and `PERUNITS=DAYS` (assuming that `DAYOFWEEK=ANY`), the period between startup windows is scheduled every 5 days after the initial start date and start time. The period between startup windows must exceed the duration of each window. Possible values are:

Hours

Specifies that time between startup windows for the schedule is specified in hours.

Days

Specifies that time between startup windows for the schedule is specified in days.

Weeks

Specifies that time between startup windows for the schedule is specified in weeks.

Months

Specifies that time between startup windows for the schedule is specified in months.

Years

Specifies that time between startup windows for the schedule is specified in years.

Onetime

Specifies that the schedule executes only once. This value overrides the value specified by the PERIOD parameter.

DAYofweek=dayofweekvalue

Specifies the day of the week on which the startup window begins. If the beginning of the calculated startup window falls on a day that does not correspond with your value for the DAYOFWEEK parameter, the startup window will be shifted forward in 24-hour increments until the DAYOFWEEK parameter is satisfied.

If you select a value for DAYOFWEEK other than ANY, depending on the values for PERIOD and PERUNITS, schedules may not be executed when you would expect. Use the QUERY EVENT command to project when schedules will be executed and thereby ensure that the desired result is achieved.

Possible values are:

ANY

Specifies that the startup window can begin on any day of the week.

WEEKDay

Specifies that the startup window can begin on Monday, Tuesday, Wednesday, Thursday, or Friday.

WEEKEnd

Specifies that the startup window can begin on Saturday or Sunday.

SUNday

Specifies that the startup window begins on Sunday.

Monday

Specifies that the startup window begins on Monday.

TUesday

Specifies that the startup window begins on Tuesday.

Wednesday

Specifies that the startup window begins on Wednesday.

THursday

Specifies that the startup window begins on Thursday.

Friday

Specifies that the startup window begins on Friday.

SAaturday

Specifies that the startup window begins on Saturday.

EXPIration=*expirationdate*

Specifies the date on which this schedule will no longer be used, in MM/DD/YYYY format. This parameter is optional. Possible values are:

Never

Specifies that the schedule never expires.

expirationdate

Specifies the date on which the schedule will no longer be used, in MM/DD/YYYY format. If you specify an expiration date, the schedule expires at 23:59:59 on the specified date.

Updating Administrative Command Schedules

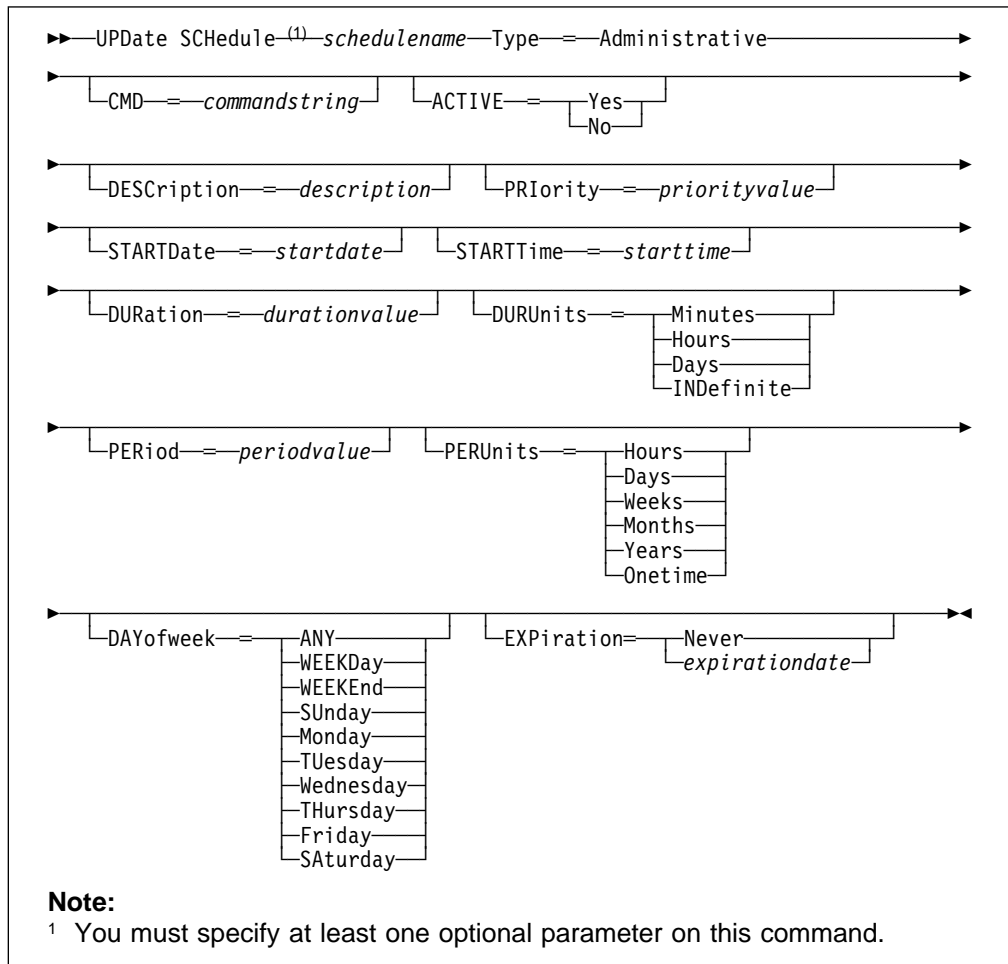
Use the UPDATE SCHEDULE command to modify selected parameters for the specified schedule. If you do not explicitly update a parameter, it remains unchanged.

For each schedule, a startup window is specified. The startup window is the time period during which the schedule must be initiated. The schedule will not necessarily complete execution within this window. If the server is not running when this window starts but is started before the end of the defined window is reached, the schedule will run when the server is restarted.

Privilege Class

To issue this command, you must have system privilege.

Syntax for Updating Administrative Command Schedules



Parameters for Updating Administrative Command Schedules

schedulename

Specifies the name of the schedule to be updated. This parameter is required.

Type=Administrative

Specifies that the schedule to be updated is an administrative schedule. This parameter is required for updating administrative command schedules.

CMD=commandstring

Specifies the actual administrative command that is to be scheduled. This parameter is optional. The maximum length of the command is 512 characters. The administrative command must be enclosed in quotation marks if it contains any blank characters.

ACTIVE=activevalue

Specifies whether the defined administrative command is eligible for scheduled execution. This parameter is optional. A defined administrative command schedule will not be executed unless it is set into the active state. Possible values are:

YES

Specifies that the defined administrative command is eligible for scheduled execution.

NO Specifies that the defined administrative command is not eligible for scheduled execution.

DESCription=*description*

Specifies a description of the schedule. This parameter is optional. The maximum length of the description is 255 characters. The description must be enclosed in quotation marks if it contains any blank characters. To remove a previously defined description, specify a null string ("") for this value.

PRiority=*priorityvalue*

Specifies the priority value for a schedule. This parameter is optional. You can specify an integer from 1–10, inclusive, with 1 being the highest priority and 10 being the lowest.

If two or more schedules for a given client have the same window start time, this value determines the order of execution. The schedule with the highest priority starts first. For example, a schedule with PRIORITY=3 starts before a schedule with PRIORITY=5.

STARTDate=*startdate*

Specifies the date for the beginning of the window in which the schedule is first executed. This parameter is optional. This parameter is used in conjunction with the STARTTIME parameter to specify when the initial startup window begins.

Possible values are:

MM/DD/YYYY

Specifies the date for the beginning of the startup window in MM/DD/YYYY format (for example, 01/31/1995).

TODAY

Specifies that the beginning of the startup window is on the current date.

TODAY+*numdays* or +*numdays*

Specifies that the beginning of the startup window is a specific number of days after the current date. The maximum value for the number of days is 9999.

For example, to specify that the beginning of the initial startup window is 3 days from now, you can specify STARTDATE=TODAY+3 or STARTDATE=+3.

STARTTime=*starttime*

Specifies the time for the beginning of the window in which the schedule is first executed, in HH:MM format. This parameter is optional. This parameter is used in conjunction with the STARTDATE parameter to specify when the startup window begins.

DURation=*durationvalue*

Specifies the number of units that define the length of the startup window. This parameter is optional. You can specify an integer in the range of 1–999, inclusive.

UPDATE SCHEDULE

This parameter is used in conjunction with the DURUNITS parameter to specify the length of the startup window. For example, if DURATION=20 and DURUNITS=MINUTES, the schedule must be started within 20 minutes of the start date and start time. The duration of the window must be shorter than the period between windows.

This value is ignored if you specify DURUNITS=INDEFINITE.

DURUnits=*durunitsvalue*

Specifies the time units used to determine the duration of the window in which this schedule can start. This parameter is optional.

This parameter is used in conjunction with the DURATION parameter to specify how long the startup window remains open to execute the schedule. For example, if DURATION=20 and DURUNITS=MINUTES, the schedule must be started within 20 minutes of the start date and start time. The duration of the window must be shorter than the period between windows. Possible values are:

Minutes

Specifies that the duration of the window is defined in minutes.

Hours

Specifies that the duration of the window is defined in hours.

Days

Specifies that the duration of the window is defined in days.

INDefinite

Specifies that the startup window has an indefinite duration. The schedule can run any time after the scheduled start time, until the schedule expires. You cannot specify DURUNITS=INDEFINITE unless PERUNITS=ONETIME is also specified for the schedule.

PERiod=*periodvalue*

Specifies the length of time between startup windows for this schedule. This parameter is optional. You can specify an integer in the range of 1–999, inclusive.

This parameter is used in conjunction with the PERUNITS parameter to specify the period between startup windows. For example, if you specify PERIOD=5 and PERUNITS=DAYS (assuming that DAYOFWEEK=ANY), the operation is scheduled every 5 days after the initial start date and start time. The period between startup windows must exceed the duration of each window.

This value is ignored if you specify PERUNITS=ONETIME.

PERUnits=*perunitsvalue*

Specifies the time units used to determine the period between startup windows for this schedule. This parameter is optional.

This parameter is used in conjunction with the PERIOD parameter to specify the period between startup windows. For example, if PERIOD=5 and PERUNITS=DAYS (assuming that DAYOFWEEK=ANY), the period between startup windows is scheduled every 5 days after the initial start date and start time. The period between startup windows must exceed the duration of each window. Possible values are:

Hours

Specifies that time between startup windows for the schedule is specified in hours.

Days

Specifies that time between startup windows for the schedule is specified in days.

Weeks

Specifies that time between startup windows for the schedule is specified in weeks.

Months

Specifies that time between startup windows for the schedule is specified in months.

Years

Specifies that time between startup windows for the schedule is specified in years.

Onetime

Specifies that the schedule executes only once. This value overrides the value specified by the PERIOD parameter.

DAYofweek=*dayofweekvalue*

Specifies the day of the week on which the startup window begins. If the beginning of the calculated startup window falls on a day that does not correspond with your value for the DAYOFWEEK parameter, the startup window will be shifted forward in 24-hour increments until the DAYOFWEEK parameter is satisfied.

If you select a value for DAYOFWEEK other than ANY, depending on the values for PERIOD and PERUNITS, schedules may not be executed when you would expect. Use the QUERY EVENT command to project when schedules will be executed and thereby ensure that the desired result is achieved. Possible values are:

ANY

Specifies that the startup window can begin on any day of the week.

WEEKDay

Specifies that the startup window can begin on Monday, Tuesday, Wednesday, Thursday, or Friday.

WEEKEnd

Specifies that the startup window can begin on Saturday or Sunday.

SUday

Specifies that the startup window begins on Sunday.

Monday

Specifies that the startup window begins on Monday.

TUesday

Specifies that the startup window begins on Tuesday.

Wednesday

Specifies that the startup window begins on Wednesday.

UPDATE SCHEDULE

Thursday

Specifies that the startup window begins on Thursday.

Friday

Specifies that the startup window begins on Friday.

Saturday

Specifies that the startup window begins on Saturday.

EXPIRATION=*expirationdate*

Specifies the date on which this schedule will no longer be used, in MM/DD/YYYY format. This parameter is optional. Possible values are:

Never

Specifies that the schedule never expires.

expirationdate

Specifies the date on which the schedule will no longer be used, in MM/DD/YYYY format. If you specify an expiration date, the schedule expires at 23:59:59 on the specified date.

Examples

Task 1

Update the MONTHLY_BACKUP schedule that belongs to the STANDARD policy domain by setting its priority value to 1.

Command: UPDate SChedule standard monthly_backup PRIority=1

Output:

```
ANR2502I Schedule MONTHLY_BACKUP updated in policy domain
STANDARD.
```

Task 2

Update the WINTER schedule that belongs to the OS2DOMAIN policy domain to expire on March 5, 1995 (03/05/1995).

Command: UPDate SChedule os2domain winter EXPIration=03/05/1995

Output:

```
ANR2502I Schedule WINTER updated in policy domain OS2DOMAIN.
```

Task 3

Update existing administrative schedule named BACKUP_ARCHIVEPOOL so that starting today, the ARCHIVEPOOL primary storage pool is backed up to the RECOVERYPOOL copy storage pool every three days at 10:00 p.m.

Command: UPDate SChedule BACKUP_ARCHIVEPOOL Type=Administrative
CMD="Backup STGpool ARCHIVEPOOL RECOVERYPOOL" ACTIVE=Yes
STARTime=22:00 PERiod=3

Output:

```
ANR2532I Schedule BACKUP_ARCHIVEPOOL updated.
```

Related Commands

Table 152 provides references to related commands.

Table 152. Commands Related to UPDATE SCHEDULE

Command	Description	Reference
COPY SCHEDULE	Creates a copy of an existing schedule.	Page 77
DEFINE SCHEDULE	Defines a schedule for a client operation or an administrative command.	Page 121
DELETE SCHEDULE	Deletes a schedule from the database.	Page 176
QUERY EVENT	Displays information about scheduled and completed events for selected clients.	Page 290
QUERY SCHEDULE	Displays information about one or more schedules.	Page 335
SET MAXCMDRETRIES	Specifies the maximum number of retries after a failed attempt to execute a scheduled command.	Page 419
SET RETRYPERIOD	Specifies the time between retry attempts by the client scheduler.	Page 430

UPDATE STGPOOL (Change the Attributes of a Storage Pool)

Use the UPDATE STGPOOL command to change any parameters for an existing storage pool.

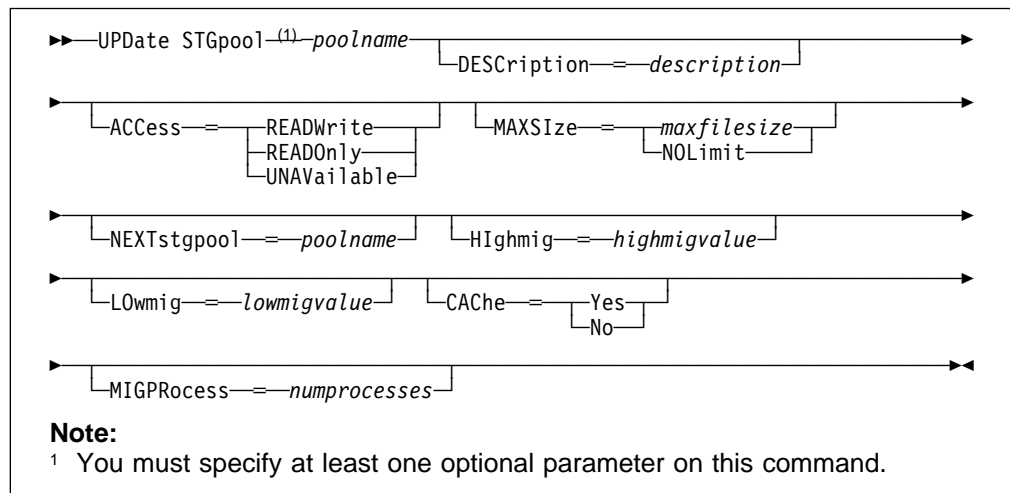
You can use this command to modify selected parameters for the specified storage pool. If you do not explicitly update a parameter, it remains unchanged.

The UPDATE STGPOOL command takes three forms, depending on whether the update is for a primary storage pool assigned to random access devices, a primary storage pool assigned to sequential access devices, or a copy storage pool. The syntax and parameters for each form are defined separately.

Privilege Class

To issue this command, you must have system privilege, unrestricted storage privilege, or restricted storage privilege for the storage pool to be updated.

Syntax for Updating a Random Access Storage Pool



Parameters for Updating a Random Access Storage Pool

poolname

Specifies the name of the storage pool to be updated. This parameter is required.

DESCRiption=*description*

Specifies a description of the storage pool to be updated. This parameter is optional. The maximum length of the description is 255 characters. The description must be enclosed in quotation marks if it contains any blank characters. To remove a previously defined description, specify a null string ("") for this value.

ACCess=*accessmode*

Specifies how users and system processes (like migration and reclamation) can access files in the storage pool. This parameter is optional. Possible values are:

READWrite

Specifies that users and system processes (like migration) can read and write to files stored on volumes in the storage pool.

READOnly

Specifies that users can only read files stored on the volumes in the storage pool.

No new writes are permitted to the volumes in the storage pool from other volumes outside the storage pool. However, system processes (like reclamation) are permitted to move files within the volumes in the storage pool.

If this storage pool has been specified as a subordinate storage pool (with the NEXTSTGPOOL parameter) and is defined as *readonly*, the storage pool is skipped when system processes (like migration) attempt to write files to the storage pool from volumes in another storage pool.

UNAVailable

Specifies that users cannot access files stored on volumes in the storage pool.

No new writes are permitted to the volumes in the storage pool from other volumes outside the storage pool. However, system processes (like reclamation) are permitted to move files within the volumes in the storage pool. Files can also be moved or copied from this storage pool to another storage pool.

If this storage pool has been specified as a subordinate storage pool (with the NEXTSTGPOOL parameter) and is defined as *unavailable*, the storage pool is skipped when system processes (like migration) attempt to write files to the storage pool.

MAXSize=*maxfilesize*

Specifies the maximum file size for the storage pool. This parameter is optional. The value of this parameter can be either NOLIMIT, or an integer followed by a scale factor. The value NOLIMIT specifies that there is no maximum size limit for files stored in the storage pool.

To limit the maximum file size, specify an integer (1–999999) followed by a scale factor. Scale factors and their meanings are:

Scale Factor	Meaning
K	Kilobyte
M	Megabyte
G	Gigabyte
T	Terabyte

For example, MAXSIZE=5G specifies that the maximum file size for this storage pool is 5 gigabytes.

If a file exceeds the maximum size and no pool is specified as the NEXTSTGPOOL, the file is not stored. If a file exceeds the maximum size and a pool is specified as the NEXTSTGPOOL, the file is stored in the next storage pool that can accept the file size. If you specify the NEXTSTGPOOL parameter, at least one storage pool in the hierarchy should have a MAXSIZE value of NOLIMIT. By setting MAXSIZE to NOLIMIT for at least one pool, you ensure that no matter what its size, a file can be stored somewhere.

NEXTstgpool=*poolname*

Specifies the name of the storage pool to which files can be migrated. This parameter is optional.

The pool specified must be a primary storage pool. A copy storage pool cannot be specified.

To remove an existing storage pool from the storage hierarchy, specify a null string ("") for this value.

If you do not specify a name for this parameter, migration cannot occur and files that exceed the value specified for the MAXSIZE parameter cannot be placed in another storage pool.

You cannot create a chain of storage pools that leads to an endless loop through the NEXTSTGPOOL parameter. At least one storage pool in the hierarchy must have no value specified for NEXTSTGPOOL.

Hlghmig=*highmigvalue*

Specifies when ADSM can start migration for this storage pool, as a percentage of the storage pool occupancy. This parameter is optional. You can specify an integer from 0–100, inclusive.

When the storage pool occupancy exceeds the value for this parameter, ADSM can start migration, by node, to the next storage pool, as defined with the NEXTSTGPOOL parameter. You can specify HIGHMIG=100 to prevent migration for this storage pool (assuming you have defined the NEXTSTGPOOL parameter).

LOWmig=*lowmigvalue*

Specifies when ADSM can stop migration for this storage pool, as a percentage of the storage pool occupancy. This parameter is optional. You can specify an integer from 0–99, inclusive.

When the storage pool occupancy reaches the value for this parameter, ADSM does not initiate migration of another node's files. Because all file spaces that belong to a node are migrated together, the occupancy of the storage pool can fall below the value you specified for the LOWMIG parameter. You can set LOWMIG=0 to permit migration to empty the storage pool.

CACHe=*cachingvalue*

For random access storage pools only, specifies whether a cached copy of a file that has been migrated remains in the random access storage pool. This parameter is optional. The primary copy of the file is migrated to the next storage pool, as defined with the NEXTSTGPOOL parameter.

Caching improves the retrievability of files that have been migrated to sequential access storage pools by permitting cached copies of files to remain in the random access storage pool. The cached copies are available, if needed, or can be overwritten when necessary. Possible values are:

Yes

Specifies that caching is enabled.

No

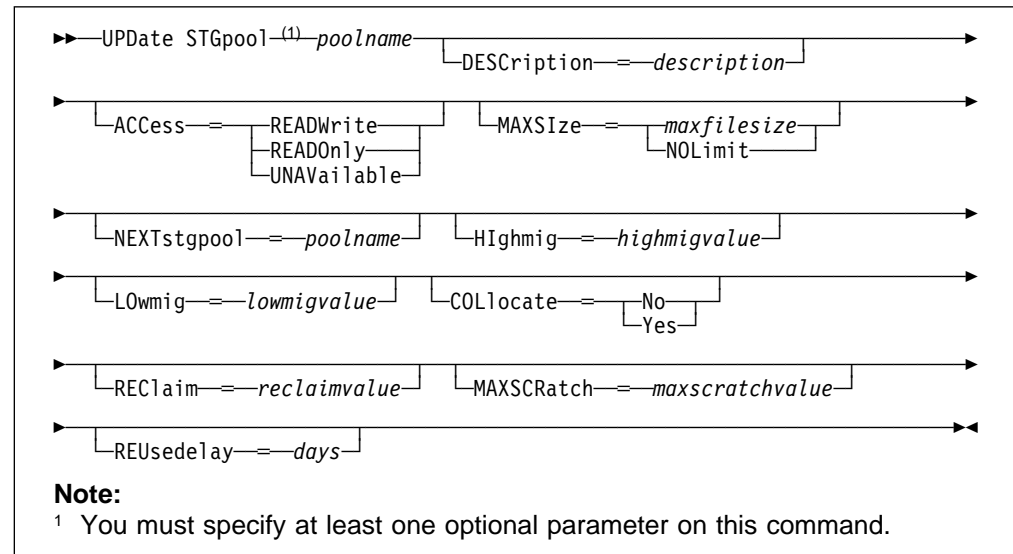
Specifies that caching is disabled.

MIGPRocess=*numprocesses*

For random access storage pools only, specifies the number of processes that are used for migrating files from this storage pool. This parameter is optional. You can specify an integer from 1–999, inclusive.

During migration, these processes are performed in parallel to provide the potential for improved migration rates.

Syntax for Updating a Primary Sequential Access Storage Pool



Parameters for Updating a Primary Sequential Access Storage Pool

poolname

Specifies the name of the storage pool to be updated. This parameter is required.

DESCRiption=*description*

Specifies a description of the storage pool to be updated. This parameter is optional. The maximum length of the description is 255 characters. The description must be enclosed in quotation marks if it contains any blank characters. To remove a previously defined description, specify a null string ("") for this value.

ACCess=*accessmode*

Specifies how users and system processes (like migration and reclamation) can access files in the storage pool. This parameter is optional. Possible values are:

READWrite

Specifies that users and system processes (like migration) can read and write files stored on volumes in the storage pool.

READOnly

Specifies that users can only read files stored on the volumes in the storage pool.

No new writes are permitted to the volumes in the storage pool from other volumes outside the storage pool. However, system processes

(like reclamation) are permitted to move files within the volumes in the storage pool.

If this storage pool has been specified as a subordinate storage pool (with the NEXTSTGPOOL parameter) and is defined as *readonly*, the storage pool is skipped when system processes (like migration) attempt to write files to the storage pool from volumes in another storage pool.

UNAVailable

Specifies that users cannot access files stored on volumes in the storage pool.

No new writes are permitted to the volumes in the storage pool from other volumes outside the storage pool. However, system processes (like reclamation) are permitted to move files within the volumes in the storage pool. Files can also be moved or copied from this storage pool to another storage pool.

If this storage pool has been specified as a subordinate storage pool (with the NEXTSTGPOOL parameter) and is defined as *unavailable*, the storage pool is skipped when system processes (like migration) attempt to write files to the storage pool.

MAXSize=*maxfilesize*

Specifies the maximum file size for the storage pool. This parameter is optional. The value of this parameter can be either NOLIMIT, or an integer followed by a scale factor. The value NOLIMIT specifies that there is no maximum size limit for files stored in the storage pool.

To limit the maximum file size, specify an integer (1–999999) followed by a scale factor. Scale factors and their meanings are:

Scale Factor	Meaning
K	Kilobyte
M	Megabyte
G	Gigabyte
T	Terabyte

For example, MAXSIZE=5G specifies that the maximum file size for this storage pool is 5 gigabytes.

If a file exceeds the maximum size and no pool is specified as the NEXTSTGPOOL, the file is not stored. If a file exceeds the maximum size and a pool is specified as the NEXTSTGPOOL, the file is stored in the next storage pool that can accept the file size. If you specify the NEXTSTGPOOL parameter, at least one storage pool in the hierarchy should have a MAXSIZE value of NOLIMIT. By setting MAXSIZE to NOLIMIT for at least one pool, you ensure that no matter what its size, a file can be stored somewhere.

NEXTstgpool=*poolname*

Specifies the name of the storage pool to which files can be migrated. This parameter is optional.

The storage pool must be a primary storage pool. A copy storage pool cannot be specified.

To remove an existing storage pool from the storage hierarchy, specify a null string ("") for this value.

If you do not specify a name for this parameter, migration cannot occur and files that exceed the value specified for the MAXSIZE parameter cannot be placed in another storage pool.

You cannot create a chain of storage pools that leads to an endless loop through the NEXTSTGPOOL parameter. At least one storage pool in the hierarchy must have no value specified for NEXTSTGPOOL.

HIGHmig=*highmigvalue*

Specifies when ADSM can start migration for this storage pool, as a percentage of the storage pool occupancy. This parameter is optional. You can specify an integer from 1–100, inclusive.

When the storage pool occupancy exceeds the value for this parameter, ADSM can start migration, by volume, to the next storage pool, as defined with the NEXTSTGPOOL parameter. You can specify HIGHMIG=100 to prevent migration for this storage pool (assuming you have defined the NEXTSTGPOOL parameter).

LOWmig=*lowmigvalue*

Specifies when ADSM can stop migration for this storage pool, as a percentage of the storage pool occupancy. This parameter is optional. You can specify an integer from 0–99, inclusive.

When the storage pool occupancy reaches the value for this parameter, ADSM will not initiate migration of files from another volume. Because all files that are on a single volume are migrated together, the occupancy of the storage pool can fall below the value you specified for the LOWMIG parameter. You can set LOWMIG=0 to permit migration to empty the storage pool.

COLlocate=*collocation*

Specifies whether ADSM attempts to keep file spaces assigned to as few volumes as possible. This parameter is optional.

Collocation reduces the number of sequential access media mounts for restore, retrieve, and recall operations. However, collocation increases both the amount of server time needed to collocate files for storing, and the number of volumes required. Collocation may not be useful on systems with only one or two sequential devices. For additional information, see *ADSM Administrator's Guide*.

Possible values are:

Yes

Specifies that collocation is enabled.

No

Specifies that collocation is disabled.

REclaim=*reclaimvalue*

Specifies when to initiate reclamation of fragmented space on volumes in the storage pool, based on the percentage of reclaimable space on a volume. Reclamation makes the fragmented space usable again by moving any remaining active files from one volume to another volume, thus making the original volume available for reuse. This parameter is optional. You can specify an integer from 1–100, inclusive.

Space becomes fragmented as files on a volume are deleted because of processes such as expiration or file space deletion. The amount of space

UPDATE STGPOOL

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 parameter, the volume is reclaimed. Active files on the volume are rewritten to other volumes in the storage pool, making the original volume available for new files.

It is recommended 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.

MAXSCRatch=*maxscratchvalue*

Specifies the maximum number of scratch volumes that can be requested by ADSM. This parameter is optional. You can specify an integer from 0–100000000, inclusive.

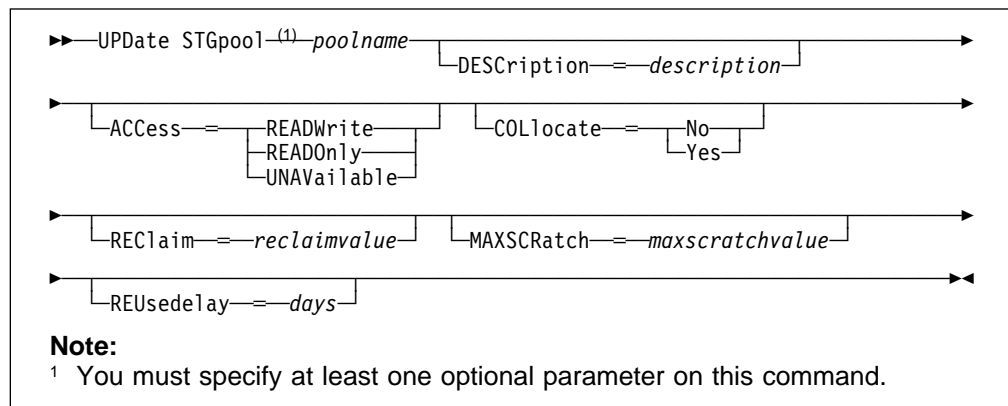
If you do not specify a value for this parameter, no scratch media are requested by ADSM.

REUsedelay=*days*

Specifies the number of days that must elapse after all files have been deleted from a volume, before the volume can be rewritten or returned to the scratch pool. This parameter is optional. You can specify an integer from 0–9999, inclusive. The value 0 means that a volume can be rewritten or returned to the scratch pool as soon as all files have been deleted from the volume.

By specifying this parameter, you can ensure that the database could be restored to an earlier level and database references to files in the storage pool would still be valid.

Syntax for Updating a Copy Storage Pool



Parameters for Updating a Copy Storage Pool

poolname

Specifies the name of the copy storage pool to be updated. This parameter is required.

DESCRiption=*description*

Specifies a description of the copy storage pool to be updated. This parameter is optional. The maximum length of the description is 255 characters. The description must be enclosed in quotation marks if it contains any blank characters. To remove a previously defined description, specify a null string (" ") for this value.

ACCess=*accessmode*

Specifies how users and system processes (like reclamation) can access files in the copy storage pool. This parameter is optional. Possible values are:

READWrite

Specifies that files can be read from and written to the volumes in the copy storage pool.

READOnly

Specifies that users can only read files stored on the volumes in the copy storage pool.

No new writes are permitted to the volumes in the copy storage pool from other volumes outside the copy storage pool. However, system processes (like reclamation) are permitted to move files within the volumes in the storage pool. Files cannot be backed up to the copy storage pool, but existing backup copies can be restored.

UNAVailable

Specifies that users cannot access files stored on volumes in the copy storage pool.

No new writes are permitted to the volumes in the copy storage pool from other volumes outside the copy storage pool. However, system processes (like reclamation) are permitted to move files within the volumes in the copy storage pool. Files cannot be backed up to the copy storage pool, but existing backup copies can be restored.

COLlocate=*collocation*

Specifies whether ADSM attempts to keep file spaces assigned to as few volumes as possible. This parameter is optional.

Collocation reduces the number of sequential access media mounts for restore, retrieve, and recall operations. However, collocation increases both the amount of server time needed to collocate files for storing, and the number of volumes required. Collocation may not be useful on systems with only one or two sequential devices. For additional information, see *ADSM Administrator's Guide*.

Possible values are:

Yes

Specifies that collocation is enabled.

No

Specifies that collocation is disabled.

REClaim=*reclaimvalue*

Specifies when to initiate reclamation of fragmented space on volumes in the copy storage pool, based on the percentage of reclaimable space on a volume. Reclamation makes the fragmented space usable again by moving any remaining active files from one volume to another volume, thus making the original volume available for reuse. This parameter is optional. You can specify an integer from 1–100, inclusive. The value 100 means that reclamation is not performed.

The amount of space that can be reclaimed on a volume increases as files on the volume are deleted because of processes such as expiration or file space deletion. It is recommended 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.

When reclamation occurs and how reclamation processing is done depends on whether the volumes are marked as offsite.

For volumes that are not marked as offsite volumes, reclamation can occur only after the volume has been filled, and then begins to empty because of file deletion. When the percentage of reclaimable space on a volume that had once been full rises above the value for the RECLAIM parameter, the volume is reclaimed. Active files on the volume are rewritten to other volumes in the storage pool, making the original volume available for new files.

For offsite volumes, reclamation can occur regardless of whether the volume has ever been filled. An offsite volume is eligible for reclamation when the percentage of unused space on the volume is greater than the reclaim parameter value. The unused space includes both space that has never been used on the volume and space that has become empty because of file deletion.

Reclamation processing also works differently for offsite storage pool volumes. When a copy pool volume that is offsite becomes eligible for reclamation, the reclamation process attempts to retrieve the active files on the reclaimable volume from a primary or copy storage pool that is onsite, and then write these files to an available volume in the original copy storage pool. Effectively, these files are moved back to the onsite location, although they could be obtained from the offsite volume after a disaster if a database backup is used which references the files on the offsite volume. Because of the way reclamation works with offsite volumes, use it carefully with copy storage pools.

MAXSCRatch=*maxscratchvalue*

Specifies the maximum number of scratch volumes that can be requested by ADSM. This parameter is optional. You can specify an integer from 0–100000000, inclusive.

If you do not specify a value for this parameter, no scratch media are requested by ADSM.

A scratch volume is automatically deleted from the copy pool after all files have been deleted and the time specified by the REUSEDELAY parameter has elapsed. However, if the access mode for a scratch volume is OFFSITE, the volume is not deleted from the copy pool until the access mode is changed; this allows an administrator to query ADSM for empty, offsite scratch volumes and return these to the onsite location.

REUsedelay=*days*

Specifies the number of days that must elapse after all files have been deleted from a volume, before the volume can be rewritten or returned to the scratch pool. This parameter is optional. You can specify an integer from 0–9999, inclusive. A value of 0 means that a volume can be rewritten or returned to the scratch pool as soon as all files have been deleted from the volume.

Important: By specifying this parameter, you can ensure that the database can be restored to an earlier level and database references to files in the copy storage pool are still valid. For this to happen, you must set this parameter to a value greater than the number of days you plan on retaining the

oldest database backup. For additional information, refer to *ADSM Administrator's Guide*.

Examples

Task 1

Update the primary storage pool named POOL1 to permit as many as 10 scratch volumes to be mounted.

Command: UPDate STGpool pool1 MAXSCRatch=10

Output:

```
ANR2202I Storage pool POOL1 updated.
```

Task 2

Update the primary storage pool named BACKUPPOOL to permit caching.

Command: UPDate STGpool BACKUPPOOL CAChe=Yes

Output:

```
ANR2202I Storage pool BACKUPPOOL updated.
```

Task 3

Update the copy storage pool named CPOOL1 to permit as many as 10 scratch volumes to be mounted and to collate files.

Command: UPDate STGpool cpool1 COLlocate=yes MAXSCRatch=10

Output:

```
ANR2202I Storage pool CPOOL1 updated.
```

Related Commands

Table 153 provides references to related commands.

Table 153 (Page 1 of 2). Commands Related to UPDATE STGPOOL

Command	Description	Reference
BACKUP DB	Backs up the ADSM database to sequential access volumes.	Page 50
BACKUP STGPOOL	Backs up a primary storage pool to a copy storage pool.	Page 56
DEFINE STGPOOL	Defines a storage pool as a named collection of data storage media.	Page 135
DELETE STGPOOL	Deletes a storage pool from data storage.	Page 179

UPDATE STGPOOL

Table 153 (Page 2 of 2). Commands Related to UPDATE STGPOOL

Command	Description	Reference
QUERY STGPOOL	Displays information about one or more storage pools.	Page 347
RESTORE STGPOOL	Restores files to a primary storage pool from copy storage pools.	Page 391
RESTORE VOLUME	Restores files stored on specified volumes in a primary storage pool from copy storage pools.	Page 396

UPDATE VOLHISTORY (Add the Location of a Database Dump, Database Backup, or Export Volume)

Use the UPDATE VOLHISTORY command to add volume location information to volume history files. You can use this location information to help track volumes that move offsite. The physical location information can be added only for a volume that contains the results of a database dump, a database backup (full or incremental), or an export operation. The location is added to the volume history files as a comment above the specified volume's detail information.

Privilege Class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```

▶▶ UPDate VOLHistory —volumename— DEVclass=—devclassname—
▶ LOcation=—location—

```

Parameters

volumename

Specifies the name of the volume for which you want to add location information. The volume must have been used for a database dump, a database backup (full or incremental), or an export operation. This parameter is required.

DEVclass=*devclassname*

Specifies the name of the device class for the volume specified. This parameter is required.

LOcation=*location*

Describes the location of the volume. This parameter is required. The maximum length of the location is 255 characters. The location must be enclosed in quotation marks if it contains any blank characters.

Examples

Task

Update the location of a volume used for database backup, BACKUP1, to show that it has been moved to an offsite location.

Command: UPDate VOLHistory backup1 DEVclass=tapebkup LOfcation="700 W. Magee Rd."

Output:

```
ANR2473I UPDATE VOLHISTORY for volume backup1 completed.
```

Related Commands

Table 154 provides references to related commands.

Table 154. Commands Related to UPDATE VOLHISTORY

Command	Description	Reference
BACKUP VOLHISTORY	Records sequential volume history information in external files.	Page 60
DELETE VOLHISTORY	Removes sequential volume history information from the volume history file.	Page 181
QUERY VOLHISTORY	Displays sequential volume history information that has been collected by the server.	Page 353

UPDATE VOLUME (Change the Attributes of One or More Storage Pool Volumes)

Use the UPDATE VOLUME command to change the access mode for one or more volumes in random or sequential access storage pools. You can also use this command to change the location information for one or more volumes in sequential access storage pools. This command can be used to update volumes assigned to primary or copy storage pools. If you do not explicitly update a parameter, it remains unchanged.

You can reset any error state associated with a volume by updating the volume to an access mode of READWRITE.

A random access volume must be varied offline before you can update the volume to *unavailable* or *destroyed*. To vary a volume offline, use the VARY command.

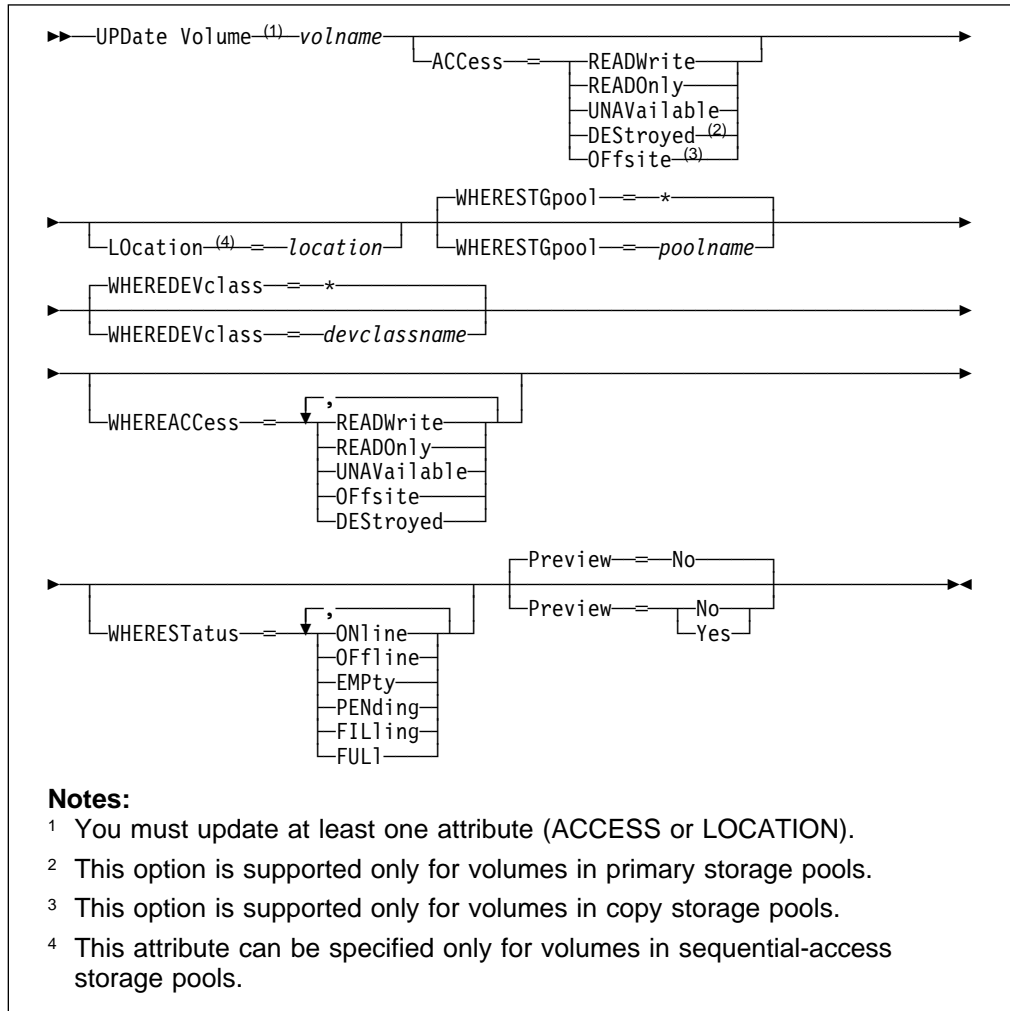
If a scratch volume with a status of EMPTY and an access mode of OFFSITE is updated so that the access mode is READWRITE, READONLY, or UNAVAILABLE, the volume is deleted from the database.

Use the QUERY VOLUME command to display the set of defined storage pool volumes.

Privilege Class

To issue this command, you must have system privilege or operator privilege.

Syntax



Parameters

volname

Specifies name of the storage pool volume to be updated. This parameter is optional. You can use a pattern-matching expression to specify this name. All matching volumes are included in the update.

ACCESS=*newaccessmode*

Specifies how users and system processes (like migration) can access files in the storage pool volume. The access mode of all updated volumes will be changed to this value. This parameter is optional. Possible values are:

READWrite

Specifies that users and system processes (like migration) can read from and write to files stored on the volume.

If the volume being updated is an empty scratch volume that had an access mode of offsite, the volume is deleted from the AD SM database.

READOnly

Specifies that users and system processes (like migration) can only read files stored on the volume.

If the volume being updated is an empty scratch volume that had an access mode of offsite, the volume is deleted from the ADSM database.

UNAVailable

Specifies that neither users nor system processes (like migration) can access files stored on the volume.

If you update a random access volume to *unavailable*, you cannot vary the volume online.

If you update a sequential access volume to *unavailable*, ADSM does not attempt to mount the volume.

If the volume being updated is an empty scratch volume that had an access mode of offsite, the volume is deleted from the ADSM database.

DESTroyed

Specifies that a primary storage pool volume has been permanently damaged. Neither users nor system processes (like migration) can access files stored on the volume.

Only volumes in primary storage pools can be updated to *destroyed*.

If you update a random access volume to *destroyed*, you cannot vary the volume online.

If you update a sequential access volume to *destroyed*, ADSM does not attempt to mount the volume.

This access mode is used to indicate an entire volume that should be restored by using the RESTORE STGPOOL command. After all files on a destroyed volume have been restored, the volume is automatically deleted from the database.

If a volume contains no files and the UPDATE VOLUME command is used to change the access mode to *destroyed*, the volume is deleted from the database.

OFFsite

Specifies that a copy storage pool volume is at an offsite location from which it cannot be mounted. Only volumes in copy storage pools can be updated to *offsite*.

Use this mode to help you track volumes that you move to offsite locations. ADSM treats offsite volumes differently than other volumes:

- Mount requests are not generated for volumes located offsite
- Data can be reclaimed or moved from offsite volumes by retrieving files from other storage pools
- Empty, offsite scratch volumes are not deleted from the copy storage pool

LOCATION=*location*

Specifies the location of the volume. This parameter is optional. It can be specified only for volumes in sequential-access storage pools. The maximum length of the location is 255 characters. The location must be enclosed in quotation marks if it contains any blank characters. To remove a previously defined description, specify the null string ("").

Note: If you specify values for both the ACCESS and LOCATION parameters but the access mode cannot be updated for a particular volume, the

location attribute is also not updated for that volume. For example, if you specify ACCESS=OFFSITE and LOCATION for primary storage pool volumes, neither the access nor location values are updated because these volumes cannot be given an access mode of OFFSITE.

WHERESTGpool=*poolname*

Specifies the name of the storage pool for volumes to be updated. This parameter can be used to restrict the update by storage pool. This parameter is optional. You can use a pattern-matching expression to specify this name. If you do not specify a storage pool name, volumes belonging to any storage pool are updated.

WHEREDEVclass=*devclassname*

Specifies the name of the device class for volumes to be updated. This parameter can be used to restrict the update by device class. This parameter is optional. You can use a pattern-matching expression to specify this name. If you do not specify a device class name, volumes with any device class are updated.

WHEREACCess=*accessmode*

Specifies the current access mode of volumes to be updated. This parameter can be used to restrict the update to volumes that currently have the specified access mode.

This parameter is optional. You can specify more than one access mode in a list by separating each mode with a comma, with no intervening spaces. If you do not specify a value for this parameter, the update is not restricted by the volumes' current access mode. Possible values are:

READWrite

Specifies that volumes with an access mode of READWRITE are updated.

READOnly

Specifies that volumes with an access mode of READONLY are updated.

UNAVailable

Specifies that volumes with an access mode of UNAVAILABLE are updated.

OFFsite

Specifies that volumes with an access mode of OFFSITE are updated.

DESTroyed

Specifies that volumes with an access mode of DESTROYED are updated.

WHERESTatus=*status*

Specifies the status of volumes to be updated. This parameter can be used to restrict the update to volumes that have a specified status. This parameter is optional. You can specify more than one status in a list by separating each status with a comma, with no intervening spaces. If you do not specify a value for this parameter, the updated volumes are not restricted by volume status. Possible values are:

ONline

Specifies that volumes with a status of ONLINE are updated.

Offline

Specifies that volumes with a status of OFFLINE are updated.

EMPTy

Specifies that volumes with a status of EMPTY are updated.

PENding

Specifies that volumes with a status of PENDING are updated. These are volumes from which all files have been deleted, but the time specified by the REUSEDELAY parameter has not elapsed.

FILLing

Specifies that volumes with a status of FILLING are updated.

FULI

Specifies that volumes with a status of FULL are updated.

Preview=*value*

Specifies whether you want to preview the update operation without actually performing it. This parameter is optional. The default value is NO. Possible values are:

No

Specifies that volumes are actually updated.

Yes

Specifies that you only want to preview the update operation. This option displays the volumes that will be updated if you actually perform the update operation.

Examples

Task 1

Update a tape volume named DSMT20 to make it unavailable to users and system processes (like migration).

Command: UPDate Volume dsmt20 ACCess=UNAVailable

Output:

```
ANR2207I Volume DSMT20 updated.
```

Task 2

Update all empty, offsite volumes in the COPYPOOL storage pool. Set the access mode to READWRITE and delete the location information for the updated volumes.

Command: UPDate Volume * ACCess=READWrite LOCation=""
WHERESTGpool=copypool WHEREACCess=offsite WHEREStatus=empty

Output:

```
ANR2207I Volume ADSMSRV0.STG0 updated.  
ANR2207I Volume ADSMSRV1.STG1 updated.  
ANR2207I Volume ADSMSRV2.STG2 updated.  
ANR2207I Volume ADSMSRV0.STG3 updated.
```

Related Commands

Table 155 provides references to related commands.

Table 155. Commands Related to UPDATE VOLUME

Command	Description	Reference
DEFINE VOLUME	Assigns a volume to be used for storage within a specified storage pool.	Page 148
DELETE VOLUME	Deletes a volume from a storage pool.	Page 185
QUERY VOLUME	Displays information about one or more storage pool volumes.	Page 359
VARY	Specifies whether a disk volume is available to the server for use.	Page 504

VALIDATE POLICYSET (Verify a Policy Set)

Use the VALIDATE POLICYSET command to verify that a policy set is complete and valid prior to activating it. The validate command examines the management class and copy group definitions in a specified policy set and reports on conditions that need to be considered if the policy set is to be activated.

The VALIDATE POLICY SET command will fail if any of the following conditions exist:

- A default management class is not defined for the policy set.
- A copy group within the policy set specifies a copy storage pool as a destination.
- A management class specifies a copy pool as the destination for space-managed files.

The following conditions result in warning messages issued during VALIDATE POLICYSET command processing:

- The destinations specified in the policy set copy group definitions do not refer to defined storage pools.

If a policy set is activated with copy group definitions that do not specify valid storage pool names, the client backup or archive operations fail.

- The active policy set in a policy domain contains management classes that are not defined in the policy set being validated.

If the policy set is activated, the files bound to the deleted management class are rebound to the default management class in the new active policy set.

- The active policy set in the policy domain contains copy groups that are not defined in the policy set being activated.

Files bound to the management classes with deleted copy group definitions are no longer archived or backed up.

- The default management class does not contain a backup or archive copy group in the policy set being activated. If this default management class is used files are not backed up or archived by default.
- A storage pool specified as a destination for files migrated from client nodes (MIGDESTINATION) is not defined.
- A management class specifies that a backup version must exist before a file can be migrated from a client node (MIGREQUIRESBKUP=Yes), but the management class does not contain a backup copy group.

After you are satisfied with the results of validation processing, use the ACTIVATE POLICYSET command to permit the attributes of the validated policy set to be defined in your storage hierarchy.

VALIDATE POLICYSET

Privilege Class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the policy set belongs.

Syntax

```
►►—VALidate POLicyset—domainname—setname—◄◄
```

Parameters

domainname

Specifies the name of the policy domain to which the policy set is assigned. This parameter is required.

setname

Specifies the name of the policy set to be verified. This parameter is required.

Examples

Task 1

Validate the policy set VACATION located in the EMPLOYEE_RECORDS policy domain.

Command: VALidate POLicyset employee_records vacation

Output:

```
ANR1515I Policy set VACATION validated in
domain EMPLOYEE_RECORDS (ready for activation).
```

Task 2

Validate the policy set SET1 located in the PROG1 policy domain.

Command: VALidate POLicyset prog1 tomc

Output:

```
ANR1553W DEFAULT Management class TOM in policy set SET1
PROG1 does not have a BACKUP copygroup: files will not be
backed up by default if this set is activated.
ANR1554W DEFAULT Management class TOM in policy set SET1
PROG1 does not have a ARCHIVE copygroup: files will not be
archived by default if this set is activated.
ANR1515I Policy set TOM validated in domain TOM (ready for
activation).
```

Task 3

Validate the policy set STANDARD located in the STANDARD policy domain.

Command: VALidate POLicyset standard standard

Output:

Note: The following output is displayed when an invalid space management destination is specified.

```
ANR1557W The space management migration destination
in management class MCLASS2
does not refer to a defined storage pool: MCLASS2.
If this pool does not exist
when policy set STANDARD is activated, clients will fail
when using this management class to migrate space-managed
files to the server.
ANR1515I Policy set STANDARD validated in domain
STANDARD (ready for activation).
```

Related Commands

Table 156 provides references to related commands.

Table 156. Commands Related to VALIDATE POLICYSET

Command	Description	Reference
ACTIVATE POLICYSET	Validates and activates a policy set within a policy domain.	Page 38
COPY POLICYSET	Creates a copy of an existing policy set.	Page 75
DEFINE COPYGROUP	Defines a copy group for backup or archive processing within a specified management class.	Page 83
DEFINE MGMTCLASS	Defines a management class within the policy domain and policy set specified.	Page 114
DELETE POLICYSET	Deletes a policy set, including its management classes and copy groups, from a policy domain.	Page 174
QUERY POLICYSET	Displays information about one or more policy sets.	Page 329
UPDATE COPYGROUP	Changes one or more attributes of a copy group.	Page 439
UPDATE POLICYSET	Changes the description of a policy set.	Page 467

VARY (Bring a Random Access Volume Online or Offline)

Use the VARY command to make a random access volume available or unavailable to the server for use as a database, recovery log, or storage pool volume.

This command is only valid for volumes that reside on random access devices. This command is used during random access maintenance and for corrective action of a random access volume.


You cannot vary a random access volume online that is defined as *unavailable*.

Varying a database or recovery log volume online can start a background process when ADSM synchronizes the volume. Use the CANCEL PROCESS command to cancel this process. To display information on background processes, use the QUERY PROCESS command.

Privilege Class

To issue this command, you must have system privilege or operator privilege.

Syntax



Parameters

ONline

Specifies that the random access volume can be used by the server as a database, recovery log, or storage pool volume. You must specify either ONLINE or OFFLINE.

OFFline

Specifies that the random access volume *cannot* be used by the server as a database, recovery log, or storage pool volume. You must specify either OFFLINE or ONLINE.

volname

Specifies the volume identifier. This parameter is required.

- For MVS VSAM linear data sets—Specify a standard MVS VSAM cluster name. A VSAM cluster name consists of one or more qualifiers, which are delimited by periods. The length of each qualifier is from 1–8 alphanumeric characters. The maximum length of the cluster name is 42 characters. An example is ADSM.BFS.POOL2.VOLA.

Any volume name used by ADSM cannot contain any embedded blanks or equal signs.

Examples

Task

Make the storage pool volume WEST.SALES.NEW unavailable to the server.

Command: VARY OFFline west.sales.new

Output:

```
ANR2211I Vary off-line initiated for disk volume
WEST.SALES.NEW.
```

Related Commands

Table 157 provides references to related commands.

Table 157. Commands Related to VARY

Command	Description	Reference
CANCEL PROCESS	Cancels a background server process.	Page 64
DEFINE DBVOLUME	Assigns a volume to be used for database storage.	Page 97
DEFINE LOGVOLUME	Defines a disk volume to the recovery log.	Page 112
DEFINE VOLUME	Assigns a volume to be used for storage within a specified storage pool.	Page 148
DELETE DBVOLUME	Deletes a volume from the database.	Page 158
DELETE LOGVOLUME	Deletes a volume from the recovery log.	Page 170
DELETE VOLUME	Deletes a volume from a storage pool.	Page 185
QUERY DBVOLUME	Displays information about one or more database volumes.	Page 280
QUERY LOGVOLUME	Displays information about one or more log volumes.	Page 311
QUERY PROCESS	Displays information about one or all background server processes.	Page 332
QUERY VOLUME	Displays information about one or more storage pool volumes.	Page 359

VARY

Appendix A. Database Recovery Commands

The following commands are used to recover your server database if a catastrophic error occurs, and you have a backup copy of the database available. These commands are presented in alphabetical order, as follows:

DSMSERV DISPLAY DBBACKUPVOLUME

Use to display information about database backup volumes when volume history information is not available.

DSMSERV DISPLAY DBVOLUMES

Use to display information about all database volumes known to the server.

DSMSERV DISPLAY LOGVOLUMES

Use to display information about all recovery log volumes known to the server.


DSMSERV RESTORE DB

Use to restore a database volume or an entire database, either to its most current state or to a specific point in time. The command uses backup versions of the database created explicitly via the BACKUP DB command or automatically via the database backup trigger.

DSMSERV DISPLAY DBBACKUPVOLUME (Display Database Backup Volume Information)

Use the DSMSERV DISPLAY DBBACKUPVOLUME command to display information about database backup volumes when volume history information is not available.

Syntax



Parameters

DEVclass=*devclassname*

Specifies the name of the sequential access device class to use. Any device class you specify must be defined in an external file specified with the DEVCONFIG option in the server options file. This parameter is required.

If the device configuration file does not exist or has been lost, you can manually create one. See *ADSM Administrator's Guide*.

VOLumentnames=*volname*

Specifies the database backup volumes for which you want to display information. This parameter is required. You can specify more than one database backup volume by separating each volume name with a comma, with no intervening spaces.

Examples

Task

Display data about the database backup volume named ARIZONA.S0149505.DBB.

Command: DSMSERV DISPLAY DBBackupvolume DEVclass=flat
VOLUmenames=arizona.s0149505.dbb

Sample JCL:

```
//ADSMSRV0 JOB ,                                JOB00250
//          REGION=40M,TIME=1440,CLASS=A,MSGCLASS=H
//ADSMSRV0 EXEC PGM=V2D14VAN,DYNAMNBR=300,TIME=NOLIMIT,
// PARM='/DISPLAY DBBACKUPVOLUME DEVCLASS=FLAT
//          VOLUMENAMES=ARIZONA.S0149505.DBB'
//STEPLIB DD DSN=EDC.V2R1M0.SEDCLINK,DISP=SHR
//          DD DSN=PLI.MINI.SIBMLINK,DISP=SHR
//          DD DSN=SYS1.LINKLIB,DISP=SHR
//          DD DSN=DSM.LINKLIB,DISP=SHR
//OPT DD DSN=ADSMSRV0.OPTIONS,DISP=SHR,FREE=CLOSE
//DSMAMENG DD DSN=TRIBLE.V2D14.MESSAGE(ANRMENU),DISP=SHR
//HLPAMENG DD DSN=TRIBLE.V2D14.HELP(ANRHENU),DISP=SHR
//DSK DD DSN=ADSMSRV0.DISKLOG,DISP=SHR
//SYSPRINT DD SYSOUT=*
//SYSTEM DD SYSOUT=*
```

Output:

```
      Date/Time: 05/10/1995 16:45:03
      Volume Type: BACKUPFULL
      Backup Series: 1
      Backup Operation: 0
      Volume Seq: 1
      Device Class: FLAT
      Volume Name: ARIZONA.S0149505.DBB
```

where:

Date/Time

The date and time that backup data was recorded on the volume.

Volume Type

The type of backup data on the volume, that is, whether the backup recorded on the volume is a full or an incremental backup.

Backup Series

The identifying number of the backup series of which this volume is a part. Each backup series consists of a full backup and all incremental backups that apply to that full backup, up to the next full backup of the database.

Backup Operation

The operation number of this backup volume, within the backup series. The full backup within a backup series is operation 0. The first incremental backup for that full backup is operation 1, the second incremental backup is operation 2, and so on.

DSMSERV DISPLAY DBBACKUPVOLUME

Volume Seq

The sequence, or position, of the volume within the backup series. Volume sequence 1 identifies the first volume used for the first operation (a full backup), and so on. For example, if the full backup occupies three volumes, these volumes are identified as volume sequence 1, 2, and 3, respectively. The first volume of the next operation (the first incremental backup) is then volume sequence 4.

Device Class

The name of the device class that this volume is assigned to.

Volume Name

The name of the volume.

DSMSERV DISPLAY DBVOLUMES (Display Database Volume Information)

Use the DSMSERV DISPLAY DBVOLUMES command to display information about all database volumes known to the ADSM server.

Syntax

```
▶▶—DSMSERV DISPlay DBVoLumes—◀◀
```

Parameters

None.

Examples

Task

Display information about all database volumes.

Command: DSMSERV DISPlay DBVoLumes

Sample JCL:

```
//ADMSRV0 JOB ,                                JOB00250
//          REGION=40M,TIME=1440,CLASS=A,MSGCLASS=H
//ADMSRV0 EXEC PGM=V2D14VAN,DYNAMNBR=300,TIME=NOLIMIT,
// PARM='/DISPLAY DBVOLUMES'
//STEPLIB DD DSN=EDC.V2R1M0.SEDCLINK,DISP=SHR
//          DD DSN=PLI.MINI.SIBMLINK,DISP=SHR
//          DD DSN=SYS1.LINKLIB,DISP=SHR
//          DD DSN=DSM.LINKLIB,DISP=SHR
//OPT      DD DSN=ADMSRV0.OPTIONS,DISP=SHR,FREE=CLOSE
//DSMAMENG DD DSN=TRIBLE.V2D14.MESSAGE(ANRMENU),DISP=SHR
//HLPAMENG DD DSN=TRIBLE.V2D14.HELP(ANRHENU),DISP=SHR
//DSK      DD DSN=ADMSRV0.DISKLOG,DISP=SHR
//SYSPRINT DD SYSOUT=*
//SYSTEM   DD SYSOUT=*
```

Output:

```
Volume Name (Copy 1): ADMSERV.DB001
      Copy Status: Sync'd
Volume Name (Copy 2):
      Copy Status: Undefined
Volume Name (Copy 3):
      Copy Status: Undefined
Available Space (MB): 4
Allocated Space (MB): 4
      Free Space (MB): 0
```

where:

Volume Name (Copy *n*)

Specifies the name of the database volume. The output displays the status of the three possible volume copies (Copy 1, Copy 2, and Copy 3).

DSMSERV DISPLAY DBVOLUMES

Copy Status

Specifies the status of the database volume. Possible values are:

Offline

The volume is varied offline.

Stale

The volume is not synchronized with the other volumes in the database and cannot be used for storage.

Sync'd

The volume is synchronized with the other volumes in the database and can be used for storage.

Undefined

No volume has been defined.

Available Space (MB)

Specifies the total amount of space that is available, both allocated and free, for the database volume, in megabytes.

Allocated Space

Specifies the amount of space allocated to the database volume, in megabytes.

Free Space

Specifies the amount of unassigned space for this mirror set that can be allocated for use, in megabytes.

DSMSERV DISPLAY LOGVOLUMES (Display Recovery Log Volume Information)

Use the DSMSERV DISPLAY LOGVOLUMES command to display information about all recovery log volumes known to the AD SM server.

Syntax

```
▶▶—DSMSERV DISPlay LOGVoLumes—◀◀
```

Parameters

None.

Examples

Task

Display information about all recovery log volumes.

Command: DSMSERV DISPlay LOGVoLumes

Sample JCL:

```
//ADMSRV0 JOB ,                                JOB00250
//          REGION=40M,TIME=1440,CLASS=A,MSGCLASS=H
//ADMSRV0 EXEC PGM=V2D14VAN,DYNAMNBR=300,TIME=NOLIMIT,
// PARM='/DISPLAY LOGVOLUMES'
//STEPLIB DD DSN=EDC.V2R1M0.SEDCLINK,DISP=SHR
//          DD DSN=PLI.MINI.SIBMLINK,DISP=SHR
//          DD DSN=SYS1.LINKLIB,DISP=SHR
//          DD DSN=DSM.LINKLIB,DISP=SHR
//OPT      DD DSN=ADMSRV0.OPTIONS,DISP=SHR,FREE=CLOSE
//DSMAMENG DD DSN=TRIBLE.V2D14.MESSAGE(ANRMENU),DISP=SHR
//HLPAMENG DD DSN=TRIBLE.V2D14.HELP(ANRHENU),DISP=SHR
//DSK      DD DSN=ADMSRV0.DISKLOG,DISP=SHR
//SYSPRINT DD SYSOUT=*
//SYSTEM   DD SYSOUT=*
```

Output:

```
Volume Name (Copy 1): ADMSERV.LOG001
      Copy Status: Sync'd
Volume Name (Copy 2):
      Copy Status: Undefined
Volume Name (Copy 3):
      Copy Status: Undefined
Available Space (MB): 8
Allocated Space (MB): 8
      Free Space (MB): 0
```

where:

Volume Name (Copy *n*)

Specifies the name of the recovery log volume. The output displays status of the three possible volume copies (Copy 1, Copy 2, and Copy 3).

DSMSERV DISPLAY LOGVOLUMES

Copy Status

Specifies the status of the recovery log volume. Possible values are:

Offline

The volume is varied offline.

Stale

The volume is not synchronized with the other volumes in the recovery log and cannot be used for storage.

Sync'd

The volume is synchronized with the other volumes in the recovery log and can be used for storage.

Undefined

No volume has been defined.

Available Space (MB)

Specifies the total amount of space that is available, both allocated and free, for the recovery log, in megabytes.

Allocated Space

Specifies the amount of space allocated for the recovery log, in megabytes.

Free Space

Specifies the amount of unassigned space in the recovery log that can be allocated for use, in megabytes.

DSMSERV RESTORE DB (Restore the Database)

Use the DSMSERV RESTORE DB command to:

- Restore an ADSM database to its most current state
- Restore a single ADSM database volume to its most current state
- Restore an ADSM database to a specific point in time when the volume history file is available
- Restore an ADSM database to a specific point in time when the volume history file is not available

The DSMSERV RESTORE DB command uses backup versions of the database you create using the BACKUP DB command or backup versions created automatically by ADSM based on the settings you specify with the DEFINE DBBACKUPTRIGGER command. You cannot use the DSMSERV RESTORE DB command to load an offline salvage dump created with the DSMSERV DUMPDB command.

To restore an ADSM database, you can use one of the following:

Roll-forward recovery

Restores the database or a single database backup volume to its most current state by first loading the most recent database backup series (the last full backup and any incremental backups that apply to that full backup), and then using recovery log records to apply any changes made to the database since the last backup was created

Point-in-time recovery

Restores the database to the point in time when a specific backup version was created

After you perform a point-in-time recovery, you must use the AUDIT VOLUME command to audit all DISK volumes and to check for and resolve any inconsistencies between database information and storage pool volumes. If you saved a copy of the volume history file before restoring the database, examine that file to find out about sequential access storage pool volumes that were deleted or reused since the point in time to which the database was restored.

The syntax and parameters for the DSMSERV RESTORE DB command vary depending on the type of database recovery you want to perform. The types are:

- Restoring a database to its most current state
- Restoring a single database volume to its most current state
- Restoring a database to a specific point in time—volume history file available
- Restoring a database to a specific point in time—volume history file unavailable

Syntax for Restoring a Database to Its Most Current State


You can perform a roll-forward recovery for an entire database, using the syntax that follows, if the following are true:

- The log mode was set to ROLLFORWARD continuously from the time that the last backup series was created, until the time that the database was damaged or lost.
- The recovery log is available and intact.
- A volume history file is available and intact.

DSMSERV RESTORE DB

When you enter the DSMSERV RESTORE DB command with no parameters, ADSM restores the database to its most current state as follows:

1. Reads the recovery log to determine which volumes contain the most recent database backup series.
2. Requests volume mounts to load the most recent backup series (the last full backup and any incremental backups that apply to that full backup). The requests for mounts and the loading continue until the entire backup series is loaded.
3. Uses the recovery log to update the database to its most current state.



Parameters for Restoring a Database to Its Most Current State

Preview=value

Specifies whether you want ADSM to display a list of the volumes needed to restore the database to its most current state, without performing the restore operation. This parameter is optional. The default value is NO. Possible values are:

No

Specifies that you want to restore the database.

Yes

Specifies that you only want to display the list of database backup volumes, without actually restoring the database.

Syntax for Restoring a Single Database Volume to Its Most Current State


You can perform a roll-forward recovery for a single database volume, using the syntax that follows, if both of the following are true:

- The log mode was set to ROLLFORWARD continuously from the time that the last backup series was created, until the time that the database volume was damaged or lost.
- The recovery log is available and intact.

When you use the DSMSERV RESTORE DB command to restore a single database volume, ADSM performs the same steps as those described in "Syntax for Restoring a Database to Its Most Current State" on page 515, except that it restores only the pages that correspond to the database volume specified.

When ADSM restores a single database volume that is mirrored, only one copy of the mirrored volumes is restored. You must use the VARY command with the ONLINE parameter to make the other mirrored volumes available again. ADSM then synchronizes the other mirrored volumes with the restored volume.

The DBVOL option applies to the current database configuration only.



Parameters for Restoring a Single Database Volume to Its Most Current State

DBVol=volname

Specifies the name of the database volume to restore. This parameter is required.

If the database volume is mirrored, you can specify the name of any volume that contains a mirrored copy of the volume you want to restore. Use the DSMSERV DISPLAY DBVOLUMES command to find the volume names.

Preview=value

Specifies whether you want ADSM to display a list of the database backup volumes needed to restore the database volume you specify, without performing the restore. This parameter is optional. The default value is NO. Possible values are:

No

Specifies that you want to restore the database volume.

Yes

Specifies that you only want to display the list of database backup volumes needed, without actually restoring the database volume.

Syntax for Restoring a Database to a Specific Point in Time—History File Available


If a backup copy of volume history information is available, you can restore a database to a specific point in time using the syntax defined below. At least one volume history file must be defined with the VOLUMEHISTORY option in the server options file, and at least one device configuration file must be defined with the DEVCONFIG option in the server options file. The server program must be able to open and read the files successfully when running the DSMSERV RESTORE DB command.

If the volume history file is not available, and you want to perform a point-in-time recovery, you can use the recovery method described in “Syntax for Restoring a Database to a Specific Point in Time—History File Unavailable” on page 519.

If a device configuration file is not available, you can manually create one. See *ADSM Administrator's Guide*.

ADSM restores the database as follows:

1. Uses the volume history file to locate the database backup volumes needed
2. Requests mounts and loads data from database backup volumes as required to restore the database volume to the specified time



Parameters for Restoring a Database to a Specific Point in Time—History File Available

TODate=*date*

Specifies the date to which to restore the database. This parameter is required. Possible values are:

MM/DD/YYYY

Specifies that you want to restore a database using the most recent backup series that was created prior to a specific date.

TODAY

Specifies that you want to restore a database using the most recent backup series that was created prior to the current date.

TODAY-*numdays* or -*numdays*

Specifies that you want to restore a database using the most recent backup series that was created prior to a specific number of days before the current date.

For example, to restore a database using the most recent backup series that was created prior to one week ago, you can enter either of the following:

```

dsmserv restore db todate=today-7
dsmserv restore db todate=-7
  
```

TOTime=*time*

Specifies the time of day to which to restore the database. This parameter is optional. The default is the end of the day (23:59:59). Possible values are:

HH:MM:SS

Specifies that you want to restore the database using the last backup series created on or before a specific time on the specified date.

NOW

Specifies that you want to restore the database using a backup series created on or before the current time on the specified date.

For example, if you issue the DSMSERV RESTORE DB command at 9:00 with TOTIME=NOW, ADSM restores the database using the last backup series created on or prior to 9:00 on the specified date.

NOW-*numhours:numminutes* or -*numhours:numminutes*

Specifies that you want to restore the database using a backup series created on or before the current time minus a specified number of hours (and optionally, minutes) on the specified date.

For example, if you issue the DSMSERV RESTORE DB command at 9:00 with TOTIME=NOW-3:30 or TOTIME=-3:30, ADSM restores the

database using the last backup series created on or prior to 5:30 on the specified date.

The following are sample commands to perform that task:

```
dsmserv restore db todate=today-7 totime=now-3:30
dsmserv restore db todate=-7 totime=-3:30
```

Preview=*value*

Specifies whether you want ADSM to display a list of the volumes needed to restore the database to the date and time you specify, without actually performing the restore operation. This parameter is optional. The default value is NO. Possible values are:

No

Specifies that you want to perform the restore operation.

Yes

Specifies that you only want ADSM to display the list of volumes needed to restore the database to the date and time you specify, without actually performing the restore operation.

Syntax for Restoring a Database to a Specific Point in Time—History File Unavailable


If the volume history file is unavailable, you can use one or more DSMSERV RESTORE DB commands with the syntax shown below to restore an ADSM database to a specific point in time.

For example, if you need to load a full backup *and* one or more incremental backups, you can issue a DSMSERV RESTORE DB command to restore the full backup, then issue an additional DSMSERV RESTORE DB command for each incremental backup.

When you use multiple DSMSERV RESTORE DB commands, you must specify COMMIT=NO for each command except the last one you issue. For the last DSMSERV RESTORE DB command, you must specify COMMIT=YES to place the database in a consistent and usable state.

If you issue a DSMSERV RESTORE DB command with a COMMIT=NO parameter, and you decide not to load additional incremental backups, you must reissue the last DSMSERV RESTORE DB command you issued and specify COMMIT=YES.

If you issue a DSMSERV RESTORE DB command with a COMMIT=YES parameter, and you decide that you need to load one or more additional incremental backups, you must reissue all of the DSMSERV RESTORE DB commands in order, starting with the restore of the full backup and specifying COMMIT=NO for each command. Then, you can enter one or more additional DSMSERV RESTORE DB commands, and specify COMMIT=YES for the last one.



Parameters for Restoring a Database to a Specific Point in Time—History File Unavailable

DEVclass=*devclassname*

Specifies the name of the sequential access device class to use. Any device class you specify for a restore must be defined in an external device configuration file specified with the DEVCONFIG option in the server options file. This parameter is required.

If the device configuration file does not exist or has been lost, you can manually create one. See *ADSM Administrator's Guide*.

VOLUMenames=*volname*

Specifies the backup volumes to use to restore the database. This parameter is required. Possible values are:

volname

Specifies the names of one or more volumes to use to restore the database. You can specify more than one backup volume by separating each volume name with a comma, with no intervening spaces. You must list the volumes in the order in which the backup was stored.

For example, if you want to use the backup volumes TAPE01, TAPE02, and TAPE03, specify the following:

```
VOLUMENAMES=TAPE01,TAPE02,TAPE03
```

FILE:*filename*

Specifies the name of a file that contains a list of the volumes to use to restore the database. In the file, each volume name must be on a separate line. You must list the volumes in the order in which the backup was stored.

For example, if you want to use the backup volumes TAPE01, TAPE02, and TAPE03, create a file named ADSM.VOLLIST that contains these lines:

```
TAPE01
TAPE02
TAPE03
```

You can then specify the volumes for the restore command as follows:

```
VOLUMENAMES=FILE:VOLLIST
```

The volume list can also be specified as an in-stream data set by specifying the data set's DD name with this parameter, and including the data set in the JCL. You must include DD: when you specify this parameter, for example:

```
VOLUMENAMES=FILE:DD:VOLLIST
```

COMMIT=*commitvalue*

Specifies whether this is the last restore command you intend to use to restore the database. This parameter is optional. The default value is NO. Possible values are:

No

Specifies that you intend to issue one or more additional DSMSERV RESTORE DB commands that list additional volumes. The database

remains in an inconsistent and unusable state until you issue a DSMSERV RESTORE DB command with a COMMIT=YES parameter.

Specify COMMIT=NO if you plan to load additional backup volumes that contain one or more incremental backups.

Yes

Specifies that this is the last restore command you intend to use to restore the database. With this command, you are specifying the last backup volumes you want to use for the restore operation.

Specify COMMIT=YES only when you want to restore the database to the end of the last volume specified, and you do not plan to restore any additional backup volumes that contain incremental backups.

Examples

Task 1

Restore the ADSM database to its most current state. The recovery log is available and intact.

Command: DSMSERV RESTORE DB

Sample JCL:

```
//ADSMSRV0 JOB ,                                JOB00250
//          REGION=40M,TIME=1440,CLASS=A,MSGCLASS=H
//ADSMSRV0 EXEC PGM=V2D14VAN,DYNAMNBR=300,TIME=NOLIMIT,
//  PARM='/RESTORE DB'
//STEPLIB DD DSN=EDC.V2R1M0.SEDCLINK,DISP=SHR
//          DD DSN=PLI.MINI.SIBMLINK,DISP=SHR
//          DD DSN=SYS1.LINKLIB,DISP=SHR
//          DD DSN=DSM.LINKLIB,DISP=SHR
//OPT      DD DSN=ADSMSRV0.OPTIONS,DISP=SHR,FREE=CLOSE
//DSMAMENG DD DSN=TRIBLE.V2D14.MESSAGE(ANRMENU),DISP=SHR
//HLPAMENG DD DSN=TRIBLE.V2D14.HELP(ANRHENU),DISP=SHR
//DSK      DD DSN=ADSMSRV0.DISKLOG,DISP=SHR
//SYSPRINT DD SYSOUT=*
//SYSTEM   DD SYSOUT=*
```

Output:

```

ANR5962I DSMSERV generated at 14:05:28 on May  8 1995.

ADSTAR Distributed Storage Manager for MVS

ANR0900I Processing options file DD:OPT.
ANR0200I Recovery log assigned capacity is 28 megabytes.
ANR0201I Database assigned capacity is 32 megabytes.
ANR0306I Recovery log volume mount in progress.
ANR5216I FILE ARIZONA.S0237622.DBB is expected to be mounted (R/O).
ANR5216I FILE ARIZONA.S0237636.DBB is expected to be mounted (R/O).
ANR5216I FILE ARIZONA.S0238297.DBB is expected to be mounted (R/O).
ANR4620I Database backup series 1 operation 0 device class FLAT.
ANR4622I   Volume 1: ARIZONA.S0237622.DBB.
ANR4622I   Volume 2: ARIZONA.S0237636.DBB.
ANR4620I Database backup series 1 operation 1 device class FLAT.
ANR4622I   Volume 1: ARIZONA.S0238297.DBB.
ANR4636I Starting roll-forward database restore.
ANR5410I FILE volume ARIZONA.S0237622.DBB mounted.
ANR1363I Input volume ARIZONA.S0237622.DBB opened (sequence number 1).
ANR4638I Restore of backup series 1 operation 0 in progress.
ANR4639I Restored 160 of 1462 database pages.
ANR4639I Restored 336 of 1462 database pages.
ANR4639I Restored 512 of 1462 database pages.
ANR4639I Restored 688 of 1462 database pages.
ANR4639I Restored 864 of 1462 database pages.
ANR4639I Restored 1040 of 1462 database pages.
ANR4639I Restored 1216 of 1462 database pages.
ANR1365I Volume ARIZONA.S0237622.DBB closed (end reached).
ANR5209I Dismounting volume ARIZONA.S0237622.DBB (read-only access).
ANR5410I FILE volume ARIZONA.S0237636.DBB mounted.
ANR1363I Input volume ARIZONA.S0237636.DBB opened (sequence number 2).
ANR4639I Restored 1392 of 1462 database pages.
ANR1364I Input volume ARIZONA.S0237636.DBB closed.
ANR5209I Dismounting volume ARIZONA.S0237636.DBB (read-only access).
ANR4640I Restored 1462 pages from backup series 1 operation 0.
ANR5410I FILE volume ARIZONA.S0238297.DBB mounted.
ANR1363I Input volume ARIZONA.S0238297.DBB opened (sequence number 1).
ANR4638I Restore of backup series 1 operation 1 in progress.
ANR4640I Restored 319 pages from backup series 1 operation 1.
ANR4641I Sequential media log redo pass in progress.
ANR1364I Input volume ARIZONA.S0238297.DBB closed.
ANR5209I Dismounting volume ARIZONA.S0238297.DBB (read-only access).
ANR0354I Recovery log redo pass in progress.
ANR4643I Processed 4096 log records.
ANR4637I Roll-forward database restore complete.
ANR5963I ADSM server termination complete.

```

Task 2

Restore a single database volume to its most current state. The volume is ADSMSRV0.DB01. The recovery log is available and intact.

Command: DSMSERV RESTORE DB DBV01=admsrv0.db01

Sample JCL:

```
//ADMSRV0 JOB ,                                JOB00292
//          REGION=40M,TIME=1440,CLASS=A,MSGCLASS=H
//ADMSRV0 EXEC PGM=V2D14VAN,DYNAMNBR=300,TIME=NOLIMIT,
// PARM='/RESTORE DB DBV=ADMSRV0.DB01'
//STEPLIB DD DSN=EDC.V2R1M0.SEDCLINK,DISP=SHR
//          DD DSN=PLI.MINI.SIBMLINK,DISP=SHR
//          DD DSN=SYS1.LINKLIB,DISP=SHR
//          DD DSN=DSM.LINKLIB,DISP=SHR
//OPT      DD DSN=ADMSRV0.OPTIONS,DISP=SHR,FREE=CLOSE
//DSMAMENG DD DSN=TRIBLE.V2D14.MESSAGE(ANRMENU),DISP=SHR
//HLPAMENG DD DSN=TRIBLE.V2D14.HELP(ANRHENU),DISP=SHR
//DSK      DD DSN=ADMSRV0.DISKLOG,DISP=SHR
//SYSPRINT DD SYSOUT=*
//SYSTEM   DD SYSOUT=*
```

Output:

```
ANR5962I DSMSERV generated at 14:05:28 on May 8 1995.

ADSTAR Distributed Storage Manager for MVS

ANR0900I Processing options file DD:OPT.
ANR0200I Recovery log assigned capacity is 28 megabytes.
ANR0201I Database assigned capacity is 32 megabytes.
ANR0306I Recovery log volume mount in progress.
ANR5216I FILE ARIZONA.S0313229.DBB is expected to be mounted (R/O).
ANR5216I FILE ARIZONA.S0313268.DBB is expected to be mounted (R/O).
ANR5216I FILE ARIZONA.S0313908.DBB is expected to be mounted (R/O).
ANR4620I Database backup series 1 operation 0 device class FLAT.
ANR4622I Volume 1: ARIZONA.S0313229.DBB.
ANR4622I Volume 2: ARIZONA.S0313268.DBB.
ANR4620I Database backup series 1 operation 1 device class FLAT.
ANR4622I Volume 1: ARIZONA.S0313908.DBB.
ANR4636I Starting roll-forward database restore.
ANR5410I FILE volume ARIZONA.S0313229.DBB mounted.
ANR1363I Input volume ARIZONA.S0313229.DBB opened (sequence number 1).
ANR4638I Restore of backup series 1 operation 0 in progress.
ANR4639I Restored 192 of 1738 database pages.
ANR4639I Restored 400 of 1738 database pages.
ANR4639I Restored 608 of 1738 database pages.
ANR4639I Restored 816 of 1738 database pages.
ANR4639I Restored 1024 of 1738 database pages.
ANR4639I Restored 1232 of 1738 database pages.
ANR1365I Volume ARIZONA.S0313229.DBB closed (end reached).
ANR5209I Dismounting volume ARIZONA.S0313229.DBB (read-only access).
ANR5410I FILE volume ARIZONA.S0313268.DBB mounted.
ANR1363I Input volume ARIZONA.S0313268.DBB opened (sequence number 2).
ANR4639I Restored 1440 of 1738 database pages.
ANR4639I Restored 1648 of 1738 database pages.
ANR1364I Input volume ARIZONA.S0313268.DBB closed.
ANR5209I Dismounting volume ARIZONA.S0313268.DBB (read-only access).
ANR4640I Restored 1734 pages from backup series 1 operation 0.
ANR5410I FILE volume ARIZONA.S0313908.DBB mounted.
ANR1363I Input volume ARIZONA.S0313908.DBB opened (sequence number 1).
ANR4638I Restore of backup series 1 operation 1 in progress.
ANR4639I Restored 496 of 953 database pages.
ANR4640I Restored 953 pages from backup series 1 operation 1.
ANR4641I Sequential media log redo pass in progress.
ANR1364I Input volume ARIZONA.S0313908.DBB closed.
ANR5209I Dismounting volume ARIZONA.S0313908.DBB (read-only access).
ANR0354I Recovery log redo pass in progress.
ANR4637I Roll-forward database restore complete.
ANR5963I ADSM server termination complete.
```

Task 3

Restore the database to its state on May 12, 1995 at 2:25 pm. The volume history file is available.

Command: DSMSERV RESTORE DB TODate=05/12/1995 TOTIME=14:25:08

Sample JCL:

```
//ADMSRV0 JOB ,                                JOB00298
//          REGION=40M,TIME=1440,CLASS=A,MSGCLASS=H
//ADMSRV0 EXEC PGM=V2D14VAN,DYNAMNBR=300,TIME=NOLIMIT,
// PARM='/RESTORE DB TODATE=05/12/1995 TOTIME=14:25:08'
//STEPLIB DD DSN=EDC.V2R1M0.SEDCLINK,DISP=SHR
//          DD DSN=PLI.MINI.SIBMLINK,DISP=SHR
//          DD DSN=SYS1.LINKLIB,DISP=SHR
//          DD DSN=DSM.LINKLIB,DISP=SHR
//OPT DD DSN=ADMSRV0.OPTIONS,DISP=SHR,FREE=CLOSE
//DSMAMENG DD DSN=TRIBLE.V2D14.MESSAGE(ANRMENU),DISP=SHR
//HLPAMENG DD DSN=TRIBLE.V2D14.HELP(ANRHENU),DISP=SHR
//DSK DD DSN=ADMSRV0.DISKLOG,DISP=SHR
//SYSPRINT DD SYSOUT=*
//SYSTEM DD SYSOUT=*
```

Output:

```

ANR5962I DSMSERV generated at 14:05:28 on May 8 1995.

ADSTAR Distributed Storage Manager for MVS

ANR0900I Processing options file DD:OPT.
ANR0200I Recovery log assigned capacity is 28 megabytes.
ANR0201I Database assigned capacity is 32 megabytes.
ANR4600I Processing volume history file ADSMSRV0.VOLHIST.
ANR5216I FILE ARIZONA.S0313229.DBB is expected to be mounted (R/O).
ANR5216I FILE ARIZONA.S0313268.DBB is expected to be mounted (R/O).
ANR5216I FILE ARIZONA.S0313908.DBB is expected to be mounted (R/O).
ANR4620I Database backup series 1 operation 0 device class FLAT.
ANR4622I Volume 1: ARIZONA.S0313229.DBB.
ANR4622I Volume 2: ARIZONA.S0313268.DBB.
ANR4620I Database backup series 1 operation 1 device class FLAT.
ANR4622I Volume 1: ARIZONA.S0313908.DBB.
ANR4634I Starting point-in-time database restore to
date 05/12/1995 14:25:08.
ANR5410I FILE volume ARIZONA.S0313229.DBB mounted.
ANR1363I Input volume ARIZONA.S0313229.DBB opened (sequence number 1).
ANR4646I Database capacity required for restore is 32 megabytes.
ANR0300I Recovery log format started; assigned capacity 28 megabytes.
ANR4638I Restore of backup series 1 operation 0 in progress.
ANR4639I Restored 192 of 1738 database pages.
ANR4639I Restored 400 of 1738 database pages.
ANR4639I Restored 608 of 1738 database pages.
ANR4639I Restored 816 of 1738 database pages.
ANR0301I Recovery log format in progress; 4 megabytes of 28.
ANR4639I Restored 1024 of 1738 database pages.
ANR4639I Restored 1232 of 1738 database pages.
ANR1365I Volume ARIZONA.S0313229.DBB closed (end reached).
ANR5209I Dismounting volume ARIZONA.S0313229.DBB (read-only access).
ANR5410I FILE volume ARIZONA.S0313268.DBB mounted.
ANR1363I Input volume ARIZONA.S0313268.DBB opened (sequence number 2).
ANR4639I Restored 1440 of 1738 database pages.
ANR4639I Restored 1648 of 1738 database pages.
ANR0301I Recovery log format in progress; 8 megabytes of 28.
ANR0301I Recovery log format in progress; 12 megabytes of 28.
ANR0301I Recovery log format in progress; 16 megabytes of 28.
ANR0301I Recovery log format in progress; 20 megabytes of 28.
ANR0301I Recovery log format in progress; 24 megabytes of 28.
ANR0301I Recovery log format in progress; 28 megabytes of 28.
ANR0302I Recovery log formatting took 28381 milliseconds.
ANR0303I Format rate: 252.6 pages/second.
ANR0304I Page service time: 4.0 ms.
ANR0305I Recovery log format complete.
ANR1364I Input volume ARIZONA.S0313268.DBB closed.
ANR5209I Dismounting volume ARIZONA.S0313268.DBB (read-only access).
ANR4640I Restored 1738 pages from backup series 1 operation 0.
ANR5410I FILE volume ARIZONA.S0313908.DBB mounted.
ANR1363I Input volume ARIZONA.S0313908.DBB opened (sequence number 1).
ANR4638I Restore of backup series 1 operation 1 in progress.
ANR4639I Restored 496 of 953 database pages.
ANR4640I Restored 953 pages from backup series 1 operation 1.
ANR0306I Recovery log volume mount in progress.
ANR4641I Sequential media log redo pass in progress.
ANR4642I Sequential media log undo pass in progress.
ANR1364I Input volume ARIZONA.S0313908.DBB closed.
ANR5209I Dismounting volume ARIZONA.S0313908.DBB (read-only access).
ANR4644I A full backup will be required for the next database backup
ANR4644I operation.
ANR4635I Point-in-time database restore complete, restore
date 05/12/1995
ANR4635I 14:25:08.
ANR5963I ADSM server termination complete.

```

Task 4

Restore the database to the time of the most recent backup of the database. The volume history file is not available. The system administrator knows that tape volumes TAPE01, TAPE02, TAPE03, and TAPE04 contain the most recent full backup of the database.

Command: DSMSERV RESTORE DB DEVclass=CARTRIDGE VOLUmenames=FILE:DD:TAPES

Sample JCL:

```
//SERVER EXEC PGM=DSMSERV,
//      PARM='/RESTORE DB DEV=CARTRIDGE VOL=FILE:DD:TAPES'
//OPT   DD DSN=ADSM.ANRSERV.OPTIONS,DISP=SHR
//DSMAMENG DD DSN=ADSM.ANRMSG(ANRMENU),DISP=SHR
//HLPAMENG DD DSN=ADSM.ANRHLP(ANRHENU),DISP=SHR
//DSK   DD DSN=ADSM.DISKLOG,DISP=SHR
//SYSPRINT DD SYSOUT=H
//SYSTEM DD SYSOUT=H
//TAPES DD *
TAPE01
TAPE02
TAPE03
TAPE04
/*
```

Output: The output is similar to that for Task 3.

DSMSERV RESTORE DB

Appendix B. Database Salvage Commands

The following commands are used to re-establish your server database if a catastrophic error occurs, and the database cannot be recovered effectively using the DSMSERV RESTORE DB command. These commands are presented in the order in which they must be executed, as follows:

DSMSERV DUMPDB

Use to dump the damaged database onto a new database.

DSMSERV INSTALL

Use to initialize the server's database and recovery log.

DSMSERV LOADDDB

Use to reload the database.

DSMSERV AUDITDB

Use to ensure that the database is returned to a consistent state after it is reloaded.

The DSMSERV DUMPDB and DSMSERV LOADDDB commands use the device configuration file. This file contains a copy of the device class definitions that exist in your database. ADSM automatically creates and maintains a current device configuration file if you specify the DEVCONFIG option in the server options file. You can also create the file by using the BACKUP DEVCONFIG command. If no device configuration file has been maintained or created, you must manually create one with an editor. For information on how to create a device configuration file manually, see *ADSM Administrator's Guide*.

DSMSERV DUMPDB (Dump the Database)

Use the DSMSERV DUMPDB command as the first step in dumping, reinitializing, and reloading the server database if a catastrophic error occurs (recovery log corruption, for example), and the DSMSERV RESTORE DB command cannot be used. No other server activity is allowed during DSMSERV DUMPDB command processing.


After DSMSERV DUMPDB processing is complete, do the following steps:

1. Issue the DSMSERV INSTALL command to reinitialize the database and recovery log.
2. Issue the DSMSERV LOADDDB command to reload the database.
3. If necessary, issue the DSMSERV AUDITDB command to locate and correct any database inconsistencies.

Note: A message at the end of the output from the DSMSERV LOADDDB command indicates whether a DSMSERV AUDITDB command is required. If the server was quiesced when it was halted, the DSMSERV AUDITDB command is not required.

You can specify either a device class or a *dumpname* to be used by the DSMSERV DUMPDB and DSMSERV LOADDDB commands. *You must use the same syntax on the two commands.* For example, if you use a device class on the DSMSERV DUMPDB command, you must use a device class on the DSMSERV LOADDDB command.

Syntax



Parameters

DEVclass=*devclassname*

Specifies the device class to which the database information is to be written. You must specify either this parameter or *dumpname*.

You must have a device configuration file that includes the definition for this device class. You must also specify the name of that device configuration file by using the DEVCONFIG option in your server options file. The device configuration file contains a copy of the device class definitions that exist in your database. The device configuration file should be available if you had previously included a DEVCONFIG option in the server options file and then started the server. If the device configuration file has been lost or was never created, do one of the following:

- If the server can be restarted, restart the server and issue the BACKUP DEVCONFIG command.

- If the server cannot be restarted, create the device configuration file manually with an editor. For information on how to create a device configuration file manually, see *ADSM Administrator's Guide*.

dumpname

Specifies the name of the location where the dumped database data is stored. You must specify either this parameter or the DEVCLASS parameter. Use the following format.

DD:ddname

Specifies the output file where the database is to be stored. A ddname is the 1–6 character data definition name created by the JCL DD statement.

Scratch=scratchvalue

Specifies whether scratch volumes can be used for dumping the database. The default value is YES. This parameter cannot be specified with the dumpname parameter.

Yes

Specifies that scratch volumes can be used for dumping the database. If you include a list of volumes on the VOLUMENAMES parameter, then scratch volumes are used only if there is not enough space to dump the database on the volumes specified. If the device type associated with the specified device class is FILE, file names for the scratch volumes are generated based on a time stamp.

No

Specifies that scratch volumes cannot be used for dumping the database. You must include a list of volumes on the VOLUMENAMES parameter to contain all of the database data.

VOLUMENAMES=volname

Specifies one or more volumes that are to be used to dump the database to. This parameter is optional and cannot be specified with the dumpname parameter. You can specify a list of volumes. The volume names in the list are separated by commas, with no intervening spaces. This parameter must be specified if SCRATCH=NO. If you do not specify this parameter and SCRATCH=YES is specified or assumed, scratch volumes are used. ADSM does not record the use of volumes by the DSMSERV DUMPDB command in the volume history file. Therefore, you must record the volume names used and specify them in the exact same order on a future DSMSERV LOADDB command.

Usage Notes

Consider the following factors before dumping the database:

- When dumping the database data to sequential access media, you need to estimate how much media is required to hold the data. If the server cannot initialize, a reasonable size estimate is equal to the size of your existing database volumes.

If you can bring the server up, you can estimate the size by following these steps:

Step 1. Issue the QUERY DB command.

DSMSERV DUMPDB

- Step 2. From the output of the QUERY DB command, multiply the number of used pages by the page size (4096 bytes) to determine the number of bytes currently occupied by the database.
 - Step 3. Divide the product of step 2 by the capacity of a single medium to determine an estimate of the number of media required.
- Record the order in which the volumes have been mounted during dump processing. The volume order is important during recovery processing so that the volumes can be remounted in the same order. To ensure that volumes are mounted in the correct order, pre-label the volumes with information that indicates the order in which they have been mounted. For example, label tapes as DSM001, DSM002, DSM003, and so on to indicate the order by which data is stored on the tape volumes.
 - When dumping, you can use scratch volumes to ensure that there is sufficient space to store the database data. If you use scratch volumes, be careful to record the label names and sequence for each volume mounted during the dump process. The volumes used by the DSMSERV DUMPDB command are *not* recorded in the sequential volume history information.
 - The server recovery log is not accessed during the dump process. Therefore, database entries that were not yet written to the database when the dump is performed are not recorded. During recovery from a catastrophic failure, the most recent database updates may not be recoverable.
 - The DSMSERV DUMPDB command requires standard label tapes if you are dumping to tape.

Examples

Task 1

Dump the existing ADSM database to a tape named DB0001.

For this example the following steps are required:

- Step 1. Halt the server.
- Step 2. Ensure that the DEVCONFIG option has been specified in the server options file. The device configuration file specified with this option must exist and must contain the device class definition needed for the dump operation.
- Step 3. Issue the DSMSERV DUMPDB command. Sample JCL:

```
//DUMPDB JOB ,REGION=40M,TIME=1440,CLASS=A,MSGCLASS=H
//ADSMDB EXEC PGM=ANRSERV,PARM='/DUMPDB DEV=CART VOL=DB0001',
//          DYNAMNBR=300,
//          TIME=NOLIMIT
//STEPLIB DD DSN=EDC.V2R1M0.SEDCLINK,DISP=SHR
//          DD DSN=SYS1.LINKLIB,DISP=SHR
//          DD DSN=PLI.SIBMLINK,DISP=SHR
//          DD DSN=ADSM.LINKLIB,DISP=SHR
//OPT     DD DSN=ADSM.OPTIONS,DISP=SHR,FREE=CLOSE
//DSMAMENG DD DSN=ADSM.ANRMSG(ANRMENU),DISP=SHR
//HLPAMENG DD DSN=ADSM.ANRHLP(ANRHENU),DISP=SHR
//DSK     DD DSN=ADSM.DISKLOG,DISP=SHR
```

Output:

```

ANR5962I DSMSERV generated at 12:42:02 on May 18 1994.

ANR0900I Processing options file DD:OPT.
ANR0990I ADSM server restart-recovery in progress.
ANR0200I Recovery log assigned capacity is 8 megabytes.
ANR0201I Database assigned capacity is 8 megabytes.
ANR4000I DUMPDB: Database dump process started.
*IEC501A M B38,DB0001,SL,NOCOMP,DUMPDB,ADSMDB,ADSM.DMP
IEC705I TAPE ON B38,DB0001,SL,NOCOMP,DUMPDB,ADSMDB,ADSM.DMP
IEC271I MESSAGE DISPLAY 'DB0001' ON B38 ISSUED BY JOB DUMPDB
ANR1360I Output volume DB0001 opened (sequence number 1).
ANR1361I Output volume DB0001 closed.
ANR4031I DUMPDB: Copied 44 database pages.
ANR4033I DUMPDB: Copied 1 bit vectors.
ANR4034I DUMPDB: Encountered 0 bad database pages.
ANR4036I DUMPDB: Copied 202 database entries.
ANR4037I DUMPDB: 18 Kilobytes copied.
ANR4001I DUMPDB: Database dump process completed.
IEF234E K B38,DB0001,PVT,DUMPDB,ADSMDB

```

Task 2

Dump the existing ADSM database to a file in the device class named FLAT, which has a device type of FILE.

For this example the following steps are required:

- Step 1. Halt the server.
- Step 2. Ensure that the DEVCONFIG option has been specified in the server options file. The device configuration file specified with this option must exist and must contain the device class definition needed for the dump operation.
- Step 3. Issue the DSMSERV DUMPDB command. Sample JCL:

```

//ADMSRV0 JOB , JOB00264
//          REGION=40M,TIME=1440,CLASS=A,MSGCLASS=H
//ADMSRV0 EXEC PGM=V2D14VAN,DYNAMNBR=300,TIME=NOLIMIT,
// PARM='/DUMPDB DEVCLASS=FLAT'
//STEPLIB DD DSN=EDC.V2R1M0.SEDCLINK,DISP=SHR
//          DD DSN=PLI.MINI.SIBMLINK,DISP=SHR
//          DD DSN=SYS1.LINKLIB,DISP=SHR
//          DD DSN=DSM.LINKLIB,DISP=SHR
//OPT DD DSN=ADMSRV0.OPTIONS,DISP=SHR,FREE=CLOSE
//DSMAMENG DD DSN=TRIBLE.V2D14.MESSAGE(ANRMENU),DISP=SHR
//HLPAMENG DD DSN=TRIBLE.V2D14.HELP(ANRHENU),DISP=SHR
//DSK DD DSN=ADMSRV0.DISKLOG,DISP=SHR
//SYSPRINT DD SYSOUT=*
//SYSTEM DD SYSOUT=*

```

Output:

```
ANR5962I DSMSERV generated at 14:05:28 on May 8 1995.

ADSTAR Distributed Storage Manager for MVS

ANR0900I Processing options file DD:OPT.
ANR0990I ADSM server restart-recovery in progress.
ANR0200I Recovery log assigned capacity is 28 megabytes.
ANR0201I Database assigned capacity is 32 megabytes.
ANR4000I DUMPDB: Database dump process started.
ANR5410I FILE volume ARIZONA.S0303570.DLD mounted.
ANR1360I Output volume ARIZONA.S0303570.DLD opened (sequence number 1).
ANR4013I DUMPDB: Dumped 10000 database entries (cumulative).
ANR4013I DUMPDB: Dumped 20000 database entries (cumulative).
ANR4013I DUMPDB: Dumped 30000 database entries (cumulative).
ANR4013I DUMPDB: Dumped 40000 database entries (cumulative).
ANR4013I DUMPDB: Dumped 50000 database entries (cumulative).
ANR4013I DUMPDB: Dumped 60000 database entries (cumulative).
ANR4013I DUMPDB: Dumped 70000 database entries (cumulative).
ANR4013I DUMPDB: Dumped 80000 database entries (cumulative).
ANR4013I DUMPDB: Dumped 90000 database entries (cumulative).
ANR4013I DUMPDB: Dumped 100000 database entries (cumulative).
ANR4013I DUMPDB: Dumped 110000 database entries (cumulative).
ANR4013I DUMPDB: Dumped 120000 database entries (cumulative).
ANR4013I DUMPDB: Dumped 130000 database entries (cumulative).
ANR4013I DUMPDB: Dumped 140000 database entries (cumulative).
ANR1361I Output volume ARIZONA.S0303570.DLD closed.
ANR5208I Dismounting volume ARIZONA.S0303570.DLD (updated).
ANR4031I DUMPDB: Copied 1812 database pages.
ANR4033I DUMPDB: Copied 9 bit vectors.
ANR4034I DUMPDB: Encountered 0 bad database pages.
ANR4036I DUMPDB: Copied 140725 database entries.
ANR4037I DUMPDB: 4 Megabytes copied.
ANR4001I DUMPDB: Database dump process completed.
ANR5963I ADSM server termination complete.
```

Task 3

Dump the existing ADSM database to tape dataset DBDATA.

For this example the following steps are required:

Step 1. Halt the server.

Step 2. Issue the DSMSERV DUMPDB command. Sample JCL:

```

//DUMPDB JOB ,REGION=40M,TIME=1440,CLASS=A,MSGCLASS=H
//ADSMDB EXEC PGM=ANRSERV,PARM='/DUMPDB DD:DBDATA',
//          DYNAMNBR=300,
//          TIME=NOLIMIT
//STEPLIB DD DSN=EDC.V2R1M0.SEDCLINK,DISP=SHR
//          DD DSN=SYS1.LINKLIB,DISP=SHR
//          DD DSN=PLI.SIBMLINK,DISP=SHR
//          DD DSN=ADSM.LINKLIB,DISP=SHR
//OPT DD DSN=ADSM.OPTIONS,DISP=SHR,FREE=CLOSE
//DSMAMENG DD DSN=ADSM.ANRMSG(ANRMENU),DISP=SHR
//HLPAMENG DD DSN=ADSM.ANRHLP(ANRHENU),DISP=SHR
//DSK DD DSN=ADSM.DISKLOG,DISP=SHR
//* -----
//* Dump To Scratch Tapes
//* -----
//DBDATA DD UNIT=B38,DISP=OLD,LABEL=(,SL),VOL=SER=SCRATCH,
//          DSN=ADSM.DUMPDB.DATA,DCB=(RECFM=U,LRECL=32760)

```

Output:

```

*IEF233A M B38,SCRATCH,,DUMPDB,ADSMDB,ADSM.DUMPDB.DATA
ANR5962I DSMSERV generated at 12:42:02 on May 18 1994.

ANR0900I Processing options file DD:OPT.
ANR0990I ADSM server restart-recovery in progress.
ANR0200I Recovery log assigned capacity is 8 megabytes.
ANR0201I Database assigned capacity is 8 megabytes.
ANR4000I DUMPDB: Database dump process started.
IEC705I TAPE ON B38,ST0077,SL,NOCOMP,DUMPDB,ADSMDB,ADSM.DUMPDB.DATA
ANR4031I DUMPDB: Copied 44 database pages.
ANR4033I DUMPDB: Copied 1 bit vectors.
ANR4034I DUMPDB: Encountered 0 bad database pages.
ANR4036I DUMPDB: Copied 202 database entries.
ANR4037I DUMPDB: 18 Kilobytes copied.
ANR4001I DUMPDB: Database dump process completed.
ANR5963I ADSM server termination complete.
IEF234E K B38,ST0077,PVT,DUMPDB,ADSMDB

```

DSMSERV INSTALL (Install the Database and Log)

Use the DSMSERV INSTALL command to run the server's install mode to initialize the server's database and recovery log. No other server activity is allowed during DSMSERV INSTALL command processing.

Warning: This command formats the log and database files. If any data exists in the files you specify, that data is lost. Do not use this command to install an ADSM Version 2 server over a Version 1 server.

Syntax

```
▶▶ dsmserv install nlog logname ndb dbname ▶▶
```

Parameters

nlog

Specifies the number of files to be used by the recovery log. You can specify as many files as necessary for your installation.

logname

Specifies the names of the files being used by the recovery log. Before issuing this command, you must allocate and format the files by using the ANRFMT job or the ANRFMT2 job found in ADSM.SAMPLIB. For information on using these jobs, see *ADSM Administrator's Guide*. The minimum size for the log is 9MB.

ndb

Specifies the number of files to be used by the database. You can specify as many files as necessary for your installation.

dbname

Specifies the names of the files being used by the database. Before issuing this command, you must allocate and format the files by using the ANRFMT job or the ANRFMT2 job found in ADSM.SAMPLIB. For information on using these jobs, see *ADSM Administrator's Guide*. The minimum size for the database is 5MB.

Examples

Task

Initialize the database and recovery log to prepare for reloading a dumped database.

For this example the following steps are required:

- Step 1. The server must already be halted.
- Step 2. Use the ANRFMT job or the ANRFMT2 job found in ADSM.SAMPLIB to allocate and format the files to be used for the database and recovery log.
- Step 3. Issue the DSMSERV INSTALL command.

Command: DSMSERV INSTALL 1 ADSMSRV0.RLOG 1 ADSMSRV0.DB01

Sample JCL follows:

```
//DSMINST JOB ,
//          REGION=40M,TIME=1440,CLASS=A,MSGCLASS=H,NOTIFY=DSMTST0
//DFDSM EXEC PGM=V2D14VAN,DYNAMNBR=300,
// PARM='/INSTALL 1 ADSMSRV0.RLOG 1 ADSMSRV0.DB01'
//STEPLIB DD DSN=EDC.V2R1M0.SEDCLINK,DISP=SHR
//          DD DSN=PLI.MINI.SIBMLINK,DISP=SHR
//          DD DSN=SYS1.LINKLIB,DISP=SHR
//          DD DSN=DSM.LINKLIB,DISP=SHR
//OPT      DD DSN=ADSMRV0.OPTIONS,DISP=SHR
//DSMAMENG DD DSN=ADSMRV0.AMENG3,DISP=SHR
//DSK      DD DSN=ADSMRV0.DISKLOG,DISP=OLD
//ERR      DD DSN=ADSMRV0.ERRORLOG,DISP=OLD
//SYSPRINT DD SYSOUT=*
//SYSTEM   DD SYSOUT=*
```

Output:

```
ANR5962I DSMSERV generated at 14:05:28 on May 8 1995.

ANR0900I Processing options file DD:OPT.
ANR0300I Recovery log format started; assigned capacity 8 megabytes.
ANR0301I Recovery log format in progress; 4 megabytes of 8.
ANR0301I Recovery log format in progress; 8 megabytes of 8.
ANR0302I Recovery log formatting took 5867 milliseconds.
ANR0303I Format rate: 349.1 pages/second.
ANR0304I Page service time: 2.9 ms.
ANR0305I Recovery log format complete.
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.
ANR0992I ADSM server installation complete.
ANR5963I ADSM server termination complete.
```

DSMSERV LOADDB (Reload the Database)

Use the DSMSERV LOADDB command to reload a dumped ADSM database to a newly installed ADSM server. The database must have been dumped by the DSMSERV command. You must reinitialize server database and recovery log volumes for ADSM prior to issuing the DSMSERV LOADDB command. No other server activity is allowed during DSMSERV LOADDB command processing.

When recovering from a failure, it is recommended that you reinitialize your server using a *new* database and log so that the original database and log volumes remain intact in case the dump and load process needs to be repeated.

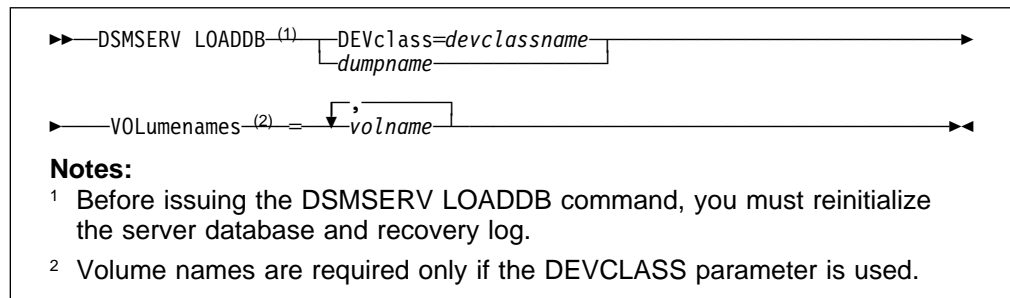
Execute a modified version of the install job (PGM=ANRSERV,PARM='/INSTALL...') to reinitialize the recovery log and database datasets. A sample job named ANRINST is provided in the samplib.

Warning: Do not reinitialize the DISK storage pool volumes.

You can specify either a device class or a *dumpname* to be used by the DSMSERV DUMPDB and DSMSERV LOADDB commands. *You must use the same syntax on the two commands.* For example, if you use a device class on the DSMSERV DUMPDB command, you must use a device class on the DSMSERV LOADDB command.

A message at the end of the output from the DSMSERV LOADDB indicates whether a DSMSERV AUDITDB command is required. If the server was quiesced when it was halted before the DSMSERV DUMPDB command was issued, the DSMSERV AUDITDB command is not required. Otherwise, you must issue the DSMSERV AUDITDB command to ensure that the database is returned to a synchronized state after it is reloaded.

Syntax



Parameters

DEVclass=classname

Specifies the device class from which the database information is to be read. You must specify either this parameter or *dumpname*.

You must have a device configuration file that includes the definition for this device class. You must also specify the name of that device configuration file by using the DEVCONFIG option in your server options file. The device configuration file contains a copy of the device class definitions that exist in your database. The device configuration file should be available if you had previously included a DEVCONFIG option in the server options file and then

started the server. If the device configuration file has been lost or was never created, and if the server cannot be restarted, create the device configuration file manually with an editor. For information on how to create a device configuration file manually, see *ADSM Administrator's Guide*.

dumpname

Specifies the name of the location where the dumped database data is stored. You must specify either this parameter or the DEVCLASS parameter. Use the following format.

DD:ddname

Specifies the file name of the dumped database data to be restored. Use the same name specified in the DSMSERV DUMPDB command. A ddname is the 1–6 character data definition name created by the JCL DD statement.

VOLUMenames=volname

Specifies one or more volumes to be used to load the database from. This parameter is required if the DEVCLASS parameter is used. The volume names in the list are separated by commas, with no intervening spaces. You must specify the same list of volumes in the same order used for the DSMSERV DUMPDB command.

Usage Notes

- Before you issue the DSMSERV LOADDB command, save your current sequential volume history to a file for future reference. The DSMSERV LOADDB process causes the sequential volume history information to be reset.
- If you must issue the DSMSERV AUDITDB command, ensure that it has completed processing before you start the server.
- If, since the dump of the database you are reloading, there have been any activities that involve the storage pool volumes, audit the volumes via the AUDIT VOLUME command. You should issue the AUDIT VOLUME command for all of your disk volumes. The volume history information that you saved earlier can help you identify which sequential volumes need to be audited. You should audit all volumes that have been reused since the time of the original dump of the database. Any volumes that were deleted from ADSM are also recorded in the sequential volume history information. If these volumes can still be obtained, you can try auditing these volumes. If these volumes cannot be obtained, you should delete references to these volumes in the database using the DELETE VOLUME command with the DISCARDATA=YES option.

Examples

Task 1

Load the ADSM database from the previously dumped DB0001 tape.

For this example the following steps are required:

- Step 1. Locate the dataset with the most recent sequential volume history information. Copy this dataset to another one, so it can be referred to later and will not be overwritten when you restart ADSM after you reload the database.
- Step 2. Re-initialize the database and log using the modified install job.

DSMSERV LOADDDB

Step 3. Ensure that the DEVCONFIG option has been specified in the server options file. The device configuration file specified with this option must exist and must contain the device class definition needed for the load operation.

Step 4. Issue the DSMSERV LOADDDB command.

Command: DSMSERV LOADDDB DEVclass=cart VOLnames=st0076

Sample JCL:

```
//LOADDB JOB ,REGION=40M,TIME=1440,CLASS=A,MSGCLASS=H
//ADSMDB EXEC PGM=ANRSERV,PARM='/LOADDB DEV=CART VOL=ST0076',
//          DYNAMNR=300,
//          TIME=NOLIMIT
//STEPLIB DD DSN=EDC.V2R1M0.SEDCLINK,DISP=SHR
//          DD DSN=SYS1.LINKLIB,DISP=SHR
//          DD DSN=PLI.SIBMLINK,DISP=SHR
//          DD DSN=ADSM.LINKLIB,DISP=SHR
//OPT     DD DSN=ADSM.OPTIONS,DISP=SHR,FREE=CLOSE
//DSMAMENG DD DSN=ADSM.ANRMSG(ANRMENU),DISP=SHR
//HLPAMENG DD DSN=ADSM.ANRHLP(ANRHENU),DISP=SHR
//DSK     DD DSN=ADSM.DISKLOG,DISP=SHR
```

Output:

```
ANR0900I Processing options file DD:OPT.

ANR0990I ADSM server restart-recovery in progress.
ANR0200I Recovery log assigned capacity is 8 megabytes.
ANR0201I Database assigned capacity is 8 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.
ANR4003I LOADDDB: Database load process started.
*IEC501A M B38,ST0076,SL,NOCOMP,LOADDB,ADSMDB,ADSM.DMP
IEC271I MESSAGE DISPLAY 'ST0076' ON B38 ISSUED BY JOB LOADDDB
ANR1363I Input volume ST0076 opened (sequence number 1).
ANR4038I LOADDDB: Loading database information dumped on 05/18/1994 at
15:15:00.
ANR4403I LOADDDB: Database dump image is consistent - a database audit
(AUDITDB)
will NOT be required if this load completes successfully.
ANR1364I Input volume ST0076 closed.
ANR4032I LOADDDB: Copied 49 database records.
ANR4033I LOADDDB: Copied 1 bit vectors.
ANR4035I LOADDDB: Encountered 0 bad database records.
ANR4036I LOADDDB: Copied 262 database entries.
ANR4037I LOADDDB: 24 Kilobytes copied.
ANR4004I LOADDDB: Database load process completed.
ANR4404I LOADDDB: Loaded a consistent dump image - a database audit
(AUDITDB) is NOT required.
IEF234E K B38,ST0076,PVT,LOADDB,ADSMDB
ANR5963I ADSM server termination complete.
```

Task 2

Load the ADSM database from the previously dumped file ARIZONA.S0303570.DLD, in the device class named FLAT.

For this example the following steps are required:

- Step 1. Locate the dataset with the most recent sequential volume history information. Copy this dataset to another one, so it can be referred to later and will not be overwritten when you restart ADSM after you reload the database.
- Step 2. Re-initialize the database and log using the modified install job.
- Step 3. Ensure that the DEVCONFIG option has been specified in the server options file. The device configuration file specified with this option must exist and must contain the device class definition needed for the load operation.
- Step 4. Issue the DSMSERV LOADDB command.

Command: DSMSERV LOADDB DEVclass=flat VOLumenames=arizona.s0303570.dld

Sample JCL:

```
//ADMSRV0 JOB ,                                JOB00274
//          REGION=40M,TIME=1440,CLASS=A,MSGCLASS=H
//ADMSRV0 EXEC PGM=V2D14VAN,DYNAMNBR=300,TIME=NOLIMIT,
//  PARM='/LOADDB DEVCLASS=FLAT VOL=ARIZONA.S0303570.DLD'
//STEPLIB DD DSN=EDC.V2R1M0.SEDCLINK,DISP=SHR
//          DD DSN=PLI.MINI.SIBMLINK,DISP=SHR
//          DD DSN=SYS1.LINKLIB,DISP=SHR
//          DD DSN=DSM.LINKLIB,DISP=SHR
//OPT      DD DSN=ADMSRV0.OPTIONS,DISP=SHR,FREE=CLOSE
//DSMAMENG DD DSN=TRIBLE.V2D14.MESSAGE(ANRMENU),DISP=SHR
//HLPAMENG DD DSN=TRIBLE.V2D14.HELP(ANRHENU),DISP=SHR
//DSK      DD DSN=ADMSRV0.DISKLOG,DISP=SHR
//SYSPRINT DD SYSOUT=*
//SYSTEM   DD SYSOUT=*
```

Output:

```

ANR5962I DSMSERV generated at 14:05:28 on May  8 1995.

ADSTAR Distributed Storage Manager for MVS

ANR0900I Processing options file DD:OPT.
ANR0990I ADSM server restart-recovery in progress.
ANR0200I Recovery log assigned capacity is 28 megabytes.
ANR0201I Database assigned capacity is 32 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.
ANR4003I LOADDB: Database load process started.
ANR5410I FILE volume ARIZONA.S0303570.DLD mounted.
ANR1363I Input volume ARIZONA.S0303570.DLD opened (sequence number 1).
ANR4038I LOADDB: Loading database information dumped on 05/12/1995 at
ANR4038I 11:32:51.
ANR4039I LOADDB: Loaded 10000 database entries (cumulative).
ANR4039I LOADDB: Loaded 20000 database entries (cumulative).
ANR4039I LOADDB: Loaded 30000 database entries (cumulative).
ANR4039I LOADDB: Loaded 40000 database entries (cumulative).
ANR4039I LOADDB: Loaded 50000 database entries (cumulative).
ANR4039I LOADDB: Loaded 60000 database entries (cumulative).
ANR4039I LOADDB: Loaded 70000 database entries (cumulative).
ANR4039I LOADDB: Loaded 80000 database entries (cumulative).
ANR4039I LOADDB: Loaded 90000 database entries (cumulative).
ANR4039I LOADDB: Loaded 100000 database entries (cumulative).
ANR4039I LOADDB: Loaded 110000 database entries (cumulative).
ANR4039I LOADDB: Loaded 120000 database entries (cumulative).
ANR4039I LOADDB: Loaded 130000 database entries (cumulative).
ANR4039I LOADDB: Loaded 140000 database entries (cumulative).
ANR1364I Input volume ARIZONA.S0303570.DLD closed.
ANR5209I Dismounting volume ARIZONA.S0303570.DLD (read-only access).
ANR4032I LOADDB: Copied 1812 database records.
ANR4033I LOADDB: Copied 9 bit vectors.
ANR4035I LOADDB: Encountered 0 bad database records.
ANR4036I LOADDB: Copied 140725 database entries.
ANR4037I LOADDB: 4 Megabytes copied.
ANR4004I LOADDB: Database load process completed.
ANR5963I ADSM server termination complete.

```

Task 3

Load the ADSM database from the previously dumped DBDATA tape dataset.

For this example the following steps are required:

- Step 1. Locate the dataset with the most recent sequential volume history information. Copy this dataset to another one, so it can be referred to later and will not be overwritten when you restart ADSM after you reload the database.
- Step 2. Re-initialize the database and log using the modified install job.
- Step 3. Issue the DSMSERV LOADDB command.

Command: DSMSERV LOADDDB DD:DBDATA

Sample JCL:

```
//LOADDB JOB ,REGION=40M,TIME=1440,CLASS=A,MSGCLASS=H
//ADSMDB EXEC PGM=ANRSERV,PARM='/LOADDB DD:DBDATA',
//          DYNAMNBR=300,
//          TIME=NOLIMIT
//STEPLIB DD DSN=EDC.V2R1M0.SEDCLINK,DISP=SHR
//          DD DSN=SYS1.LINKLIB,DISP=SHR
//          DD DSN=PLI.SIBMLINK,DISP=SHR
//          DD DSN=ADSM.LINKLIB,DISP=SHR
//OPT DD DSN=ADSM.OPTIONS,DISP=SHR,FREE=CLOSE
//DSMAMENG DD DSN=ADSM.ANRMSG(ANRMENU),DISP=SHR
//HLPAMENG DD DSN=ADSM.ANRHLP(ANRHENU),DISP=SHR
//DSK DD DSN=ADSM.DISKLOG,DISP=SHR
//* -----
//* LOAD The Database from a tape
//* -----
//DBDATA DD UNIT=B38,DISP=OLD,LABEL=(,SL),VOL=SER=ST0077,
//          DSN=ADSM.DUMPDB.DATA,DCB=(RECFM=U,LRECL=32760)
```

Output:

```
*IEF233A M B38,ST0077,,LOADDB,ADSMDB,ADSM.DUMPDB.DATA
ANR5962I DSMSERV generated at 12:42:02 on May 18 1994.

ANR0900I Processing options file DD:OPT.
ANR0990I ADSM server restart-recovery in progress.
ANR0200I Recovery log assigned capacity is 8 megabytes.
ANR0201I Database assigned capacity is 8 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.
ANR4003I LOADDDB: Database load process started.
ANR4038I LOADDDB: Loading database information dumped on 05/18/1994 at
14:54:10.
ANR4032I LOADDDB: Copied 44 database records.
ANR4033I LOADDDB: Copied 1 bit vectors.
ANR4035I LOADDDB: Encountered 0 bad database records.
ANR4036I LOADDDB: Copied 202 database entries.
ANR4037I LOADDDB: 18 Kilobytes copied.
ANR4004I LOADDDB: Database load process completed.
ANR4405I LOADDDB: Loaded an inconsistent dump image - a database audit
(AUDITDB) IS REQUIRED.
ANR5963I ADSM server termination complete.
IEF234E K B38,ST0077,PVT,LOADDB,ADSMDB
```

DSMSERV AUDITDB (Verify Database Information)


Use the DSMSERV AUDITDB command when required to ensure that the database is returned to a synchronized state after it is reloaded via the DSMSERV LOADDB command. A message at the end of the output from the DSMSERV LOADDB command indicates whether this command is required.

You must issue this command prior to starting the server. To successfully use this command, one of the following conditions must be true:

- The recovery log mode was set to normal before the server was halted (by using the SET LOGMODE command)
- A backup of the database was completed just before the server was halted

No other server activity is allowed during DSMSERV AUDITDB command processing.

Syntax



Parameters

Fix=fixvalue

Specifies how ADSM resolves any inconsistencies in the database. This parameter is optional. The default value is NO. Possible values are:

No

Specifies that ADSM reports, but does not fix, any logical inconsistencies found.

If the audit finds inconsistencies when the command is run with FIX=NO specified, before making the server available for production work you must re-issue the command specifying FIX=YES. Because AUDITDB must be run with FIX=YES to recover the database, the recommended usage in a recovery situation is FIX=YES the first time.

Yes

Specifies that ADSM fixes any inconsistencies and issues messages indicating the corrective actions.

Detail=detailvalue

Specifies the level of database testing to be performed by ADSM. This parameter is optional. The default value is NO. Possible values are:

No

Specifies that ADSM only tests referential integrity of the database.

Yes

Specifies that ADSM tests referential integrity of the database and the integrity of each individual database entry.

Examples

Task 1

Audit the ADSM database for inconsistencies, but do not fix inconsistencies if they exist.

Command: DSMSERV AUDITDB Fix=No

Sample JCL:

```
//ADMSRV0 JOB ,                                JOB00263
//          REGION=40M,TIME=1440,CLASS=A,MSGCLASS=H
//ADMSRV0 EXEC PGM=V2D14VAN,DYNAMNBR=300,TIME=NOLIMIT,
//  PARM='/AUDITDB FIX=NO'
//STEPLIB DD DSN=EDC.V2R1M0.SEDCLINK,DISP=SHR
//          DD DSN=PLI.MINI.SIBMLINK,DISP=SHR
//          DD DSN=SYS1.LINKLIB,DISP=SHR
//          DD DSN=DSM.LINKLIB,DISP=SHR
//OPT      DD DSN=ADMSRV0.OPTIONS,DISP=SHR,FREE=CLOSE
//DSMAMENG DD DSN=TRIBLE.V2D14.MESSAGE(ANRMENU),DISP=SHR
//HLPAMENG DD DSN=TRIBLE.V2D14.HELP(ANRHENU),DISP=SHR
//DSK      DD DSN=ADMSRV0.DISKLOG,DISP=SHR
//SYSPRINT DD SYSOUT=*
//SYSTEM   DD SYSOUT=*
```

Output:

```
ANR5962I DSMSERV generated at 14:05:28 on May 8 1995.

ANR0900I Processing options file DD:OPT.
ANR0990I ADSM server restart-recovery in progress.
ANR0200I Recovery log assigned capacity is 28 megabytes.
ANR0201I Database assigned capacity is 32 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.
ANR4140I AUDITDB: Database audit process started.
ANR4075I AUDITDB: Auditing policy definitions.
ANR4040I AUDITDB: Auditing client node and administrator definitions.
ANR4135I AUDITDB: Auditing central scheduler definitions.
ANR2833I AUDITDB: Auditing license definitions.
ANR4136I AUDITDB: Auditing server inventory.
ANR4306I AUDITDB: Processed 10000 database entries (cumulative).
ANR4306I AUDITDB: Processed 20000 database entries (cumulative).
ANR4137I AUDITDB: Auditing inventory file spaces.
ANR4307I AUDITDB: Auditing inventory external space-managed objects.
ANR4138I AUDITDB: Auditing inventory backup objects.
ANR4306I AUDITDB: Processed 30000 database entries (cumulative).
ANR4306I AUDITDB: Processed 40000 database entries (cumulative).
ANR4139I AUDITDB: Auditing inventory archive objects.
ANR4310I AUDITDB: Auditing inventory space-managed objects.
ANR4230I AUDITDB: Auditing data storage definitions.
ANR4264I AUDITDB: Auditing file information.
ANR4265I AUDITDB: Auditing disk file information.
ANR4306I AUDITDB: Processed 50000 database entries (cumulative).
ANR4306I AUDITDB: Processed 60000 database entries (cumulative).
ANR4306I AUDITDB: Processed 70000 database entries (cumulative).
ANR4266I AUDITDB: Auditing sequential file information.
ANR4306I AUDITDB: Processed 80000 database entries (cumulative).
ANR4306I AUDITDB: Processed 90000 database entries (cumulative).
ANR4306I AUDITDB: Processed 100000 database entries (cumulative).
ANR4271E AUDITDB: Missing or incorrect occupancy information detected.
ANR4306I AUDITDB: Processed 110000 database entries (cumulative).
ANR4306I AUDITDB: Processed 120000 database entries (cumulative).
ANR4306I AUDITDB: Processed 130000 database entries (cumulative).
ANR4256I AUDITDB: Auditing data storage definitions for disk volumes.
ANR4306I AUDITDB: Processed 140000 database entries (cumulative).
ANR4263I AUDITDB: Auditing data storage definitions for sequential
volumes.
ANR4141I AUDITDB: Database audit process completed.
ANR5963I ADSM server termination complete.
```

Task 2

Audit the ADSM database for inconsistencies and fix any inconsistencies if they exist.

Command: DSMSERV AUDITDB Fix=Yes

Output:

```
ANR7800I DSMSERV generated at 16:44:34 on Jan 25 1994.
ANR0990I ADSM server restart-recovery in progress.
ANR0200I Recovery log assigned capacity is 16 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.
ANR4140I AUDITDB: Database audit process started.
ANR4075I AUDITDB: Auditing policy definitions.
ANR4087I AUDITDB: Invalid node 88 found assigned to domain STANDARD
reference will be deleted.
ANR4077I AUDITDB: Invalid client node count detected for policy domain
STANDARD count will be corrected.
ANR4096E AUDITDB: An invalid management class identifier (22) was
encountered for management class STANDARD in policy set STANDARD,
domain STANDARD (expected 15) - the correct identifier will be stored.
ANR4090E AUDITDB: Active Policy set for domain STANDARD does not contain
a valid default management class - policy will fail for this domain.
Please correct the ACTIVE policy set.
ANR4040I AUDITDB: Auditing client node and administrator definitions.
ANR4042I AUDITDB: Node DAVE is assigned to Domain STANDARD, but not
verified by the Policy Manager - attempting to assign the node to
domain STANDARD.
ANR4066I AUDITDB: Administrator 99 not found in the server database
- reference will be removed.
ANR4135I AUDITDB: Auditing central scheduler definitions.
ANR2833I AUDITDB: Auditing license definitions.
ANR4136I AUDITDB: Auditing server inventory.
ANR4137I AUDITDB: Auditing inventory file spaces.
ANR4138I AUDITDB: Auditing inventory backup objects.
ANR4139I AUDITDB: Auditing inventory archive objects.
ANR4230I AUDITDB: Auditing data storage definitions.
ANR4264I AUDITDB: Auditing file information.
ANR4265I AUDITDB: Auditing disk file information.
ANR4266I AUDITDB: Auditing sequential file information.
ANR4256I AUDITDB: Auditing data storage definitions for disk volumes.
ANR4306I AUDITDB: 25000 Records processed.
ANR4263I AUDITDB: Auditing data storage definitions for sequential
volumes.
ANR4210I AUDITDB: Auditing physical volume repository definitions.
ANR4141I AUDITDB: Database audit process completed.
```

Glossary

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

This glossary may include terms and definitions from:

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

A

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

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

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

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

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

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

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

active version. The most recent backup copy of a file stored by ADSM. Such a file is exempt from deletion until a backup detects that the user has either replaced

the file with a newer version, or has explicitly deleted the file from the workstation. Contrast with *inactive version*.

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

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

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

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

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

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

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

ADSM. ADSTAR Distributed Storage Manager.

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

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

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

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

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

AFS. Andrew file system.

AIX. Advanced Interactive Executive.

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

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

API. Application programming interface.

APPC. Advanced Program-to-Program Communication.

APPN. Advanced Peer-to-Peer Networking.

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

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

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

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

archive retention grace period. The number of days ADSM retains an archive copy when the server is

unable to rebind the file to an appropriate management class.

AS/400. Application System/400.

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

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

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

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

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

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

B

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

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

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

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

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

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

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

backup series. See *database backup series*.

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

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

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

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

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

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

C

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

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

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

cartridge system tape (CST). The base tape cartridge media used with 3480 or 3490 Magnetic Tape Subsys-

tems. When specified as a media type in ADSM, CST identifies standard length tape. Contrast with *enhanced capacity cartridge system tape*.

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

CID. Configuration Installation and Distribution.

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

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

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

client migration. The process of copying a file from a client node to ADSM storage and replacing the file with a stub file on the client node. The process is controlled by the user and by space management attributes in the management class. See also *space management*.

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

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

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

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

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

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

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

client user options file. A user-created file containing a default set of processing options that identify the server, communication method, backup and archive options, space management options, and scheduling options. Also called the *dsm.opt* file. See also *client system options file*.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

CPI-C. Common Programming Interface Communications.

CST. Cartridge system tape.

CUA. Common User Access.

D

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

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

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

DASD. Direct access storage device.

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

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

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

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

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

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

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

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

dataserver. See *Tape Library Dataserver*.

data set. See *linear data set*.

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

DDM. Distributed Data Management.

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

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

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

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

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

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

desktop. On-screen representation of a desk top.

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

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

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

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

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

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

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

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

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

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

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

DLL. Dynamic link library.

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

DOS. Disk operating system.

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

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

dsmerv.opt. See *server options file*.

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

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

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

E

ECCST. Enhanced capacity cartridge system tape.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

F

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

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

file space. A logical space in a client's storage that can contain a group of files. For clients on systems such as OS/2, a file space is a logical partition and is identified by a volume label. For clients on systems such as AIX and UNIX, a file space can consist of any subset of directories and subdirectories stemming from a virtual mount point. Clients can restore, retrieve, or delete their file spaces from ADSM's data storage.

ADSM does not necessarily store all the files from a single file space together, but can identify all the files in its data storage that came from a single file space.

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

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

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

FTP. File Transfer Protocol.

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

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

G

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

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

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

GUI. Graphical user interface.

H

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

HDA. Head-disk assembly.

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

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

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

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

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

HSM client. Hierarchical storage management client.

I

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

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

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

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

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

include-exclude file (for UNIX clients) or in the client options file (for other clients).

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

incremental backup. (1) A function that allows users to back up files or directories from a client domain that are not excluded in the include-exclude list and that meet the requirements for frequency, mode, and serialization as defined in the backup copy group of the management class assigned to the files. Contrast with *selective backup*. (2) An ADSM function that copies only the pages in the database that are new or changed since the last full or incremental backup. Contrast with *full backup*. See *database backup series*.

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

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

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

IUCV. Inter-user communication vehicle.

K

KB. Kilobyte.

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

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

kilobyte (KB). 1024 bytes.

L

LAN. Local area network.

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

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

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

load. See *mount*.

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

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

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

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

M

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

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

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

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

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

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

MB. Megabyte.

megabyte (MB). (1) For processor storage and real and virtual memory, 2²⁰ or 1 048 576 bytes. (2) For disk storage capacity and transmission rates, 1 000 000 bytes.

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

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

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

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

mm. Millimeter.

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

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

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

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

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

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

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

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

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

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

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

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

MVS. Multiple Virtual Storage.

N

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

NETBIOS. Network Basic Input/Output System.

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

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

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

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

NFS. Network File System.

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

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

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

O

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

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

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

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

Operating System/2 (OS/2). An operating system used in IBM PC AT, PS/2, and compatible computers. OS/2 is one of the supported client environments and one of the supported server environments.

operator privilege class. An administrative privilege class that allows an administrator to issue commands that control the operation of the server. This privilege

class allows disabling or halting the server to perform maintenance, enabling the server, canceling server processes, and managing tape.

optical disk. A disk that contains data readable by optical techniques.

optical drive. A drive mechanism that rotates an optical disc.

optical library. A disk storage device that houses optical disk drives and optical disks, and contains a mechanism for moving optical disks between a storage area and optical disk drives.

OS/2. Operating System/2.

OS/400. Operating System/400.

owner. The owner of backup-archive files sent from a multiuser client node, such as AIX.

P

page. (1) A block of instructions, data, or both. (2) In ADSM, a unit of space allocation within database volumes. (3) In a virtual storage system, a fixed block that has a virtual address and is transferred as a unit between real storage and auxiliary storage.

paging. (1) The action of transferring instructions, data, or both, between real storage and external page storage. (2) Moving data between memory and a mass storage device as the data is needed.

pattern-matching expression. A string expression that uses wildcard characters to specify one or more ADSM objects. See also *wildcard character*.

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

platform. The operating system environment in which a program runs.

policy definition. Server control information that includes information about policy domains, policy sets (including the ACTIVE policy set), management classes (including the default management class), copy groups, schedules, and associations between client nodes and schedules. A policy definition can be exported from a source server so that it can be imported to a target server at a later date.

policy domain. A policy object that contains policy sets, management classes, and copy groups that is used by a group of client nodes. See *policy set*, *management class*, and *copy group*.

policy privilege class. An administrative privilege class that allows an administrator to manage policy objects, register client nodes, and schedule client operations (such as backup services) for client nodes. Administrators can be authorized with unrestricted or restricted policy privilege. See *unrestricted policy privilege* or *restricted policy privilege*.

policy set. A policy object that contains a group of management class definitions that exist for a policy domain. At any one time there can be many policy sets within a policy domain but only one policy set can be active. See *management class* and *active policy set*.

primary storage pool. A named set of volumes that ADSM uses to store backup versions of files, archive copies of files, and files migrated from client nodes via ADSM space management. A primary storage pool may be backed up to a copy storage pool either automatically or by command. See *destination* and *copy storage pool*.

privilege class. A level of authority granted to an ADSM administrator. ADSM has five privilege classes: system, policy, storage, operator, and analyst. The privilege class determines which ADSM administrative tasks the administrator can perform. For example, an administrator with system privilege class can perform any administrative task.

programmable workstation communication services (PWSCS). A product that provides transparent high performance communications between programs running on workstations or on host systems.

protection status. A device class attribute that specifies whether to update the RACF profile to identify which users have access to cartridge tapes associated with this device class on MVS servers.

PWSCS. Programmable workstation communication services.

Q

QIC. Quarter-inch cartridge (a type of magnetic tape media).

R

random access media. Any volume accessed in a nonsequential manner. In ADSM, volumes are accessed in a nonsequential manner if they reside in the DISK device class.

randomization. The percentage of the startup window that the server can use to randomize start times for different client nodes associated with a schedule.

rebinding. The process of associating a file with a new management class name. For example, rebinding occurs when the management class associated with a file is deleted. See *binding*.

recall. A function that allows users to access files that have been migrated from their workstations to ADSM storage via ADSM space management. Contrast with *migrate*.

reclamation. A process of consolidating the remaining data from many sequential access media onto a single new sequential access media.

reclamation threshold. A value that specifies a percentage of space on sequential access media volumes that can be occupied by reclaimable space. The remainder of the space is for active data. (Space becomes reclaimable when files are expired.)

recovery log. A log of updates that are about to be written to the database. The log can be used to recover from system and media failures.

recovery log buffer pool. Used to hold new transactions records until they can be written to the recovery log.

reduce. The process of freeing up enough space to allow you to delete a volume from the database or recovery log. Contrast with *extend*.

REEL. On a VM server, a device class that is defined by ADSM at installation. It is used with VM servers to categorize tape devices that support tape reels, such as the 3420 9-track tape device.

register. Defines a client node or administrator who can access the server. See *registration*.

registration. The process of identifying a client node or administrator to the server.

reply operator. On a VM server, a VM user ID that specifies an operator who will reply to tape mount requests by the server.

restore. The process of returning a backup copy to an active storage location for use. ADSM has processes for restoring its database, storage pools, storage pool volumes, and users' backed-up files. For example, users can copy a backup version of a file from the storage pool to the workstation. The backup version in the storage pool is not affected. Contrast with *backup*.

restricted policy privilege. An administrative privilege class that enables an administrator to manage policy objects only for the policy domains for which the administrator has been authorized.

restricted storage privilege. An administrative privilege class that enables an administrator to control the allocation and use of storage resources only for the storage pools for which the administrator has been authorized.

retention. The amount of time, in days, that inactive backed up or archived files will be retained in the storage pool before they are deleted. The following copy group attributes define retention: retain extra versions, retain only version, retain version.

retention period. On an MVS server, a device class attribute that specifies how long files are retained on sequential access media. When used, ADSM passes this information to the MVS operating system to ensure that other tape management systems do not overwrite tape volumes that contain retained data.

retrieve. A function that allows users to copy an archive copy from the storage pool to the workstation. The archive copy in the storage pool is not affected. Contrast with *archive*.

RLIO. Record Level Input/Output.

rollback. To remove changes that were made to database files since the last commit point.

root. In the AIX and UNIX environments, the user name for the system user with the most authority.

root user. In the AIX and UNIX environments, an expert user who can log in and execute restricted commands, shut down the system, and edit or delete protected files. Also called the *superuser*.

routing choice. A choice in a pull-down menu that, when selected, brings up another window. See also *action choice*.

S

SAA. Systems Application Architecture.

schedule. A database record that describes scheduled client operations or administrative commands. See *administrative command schedule* and *client schedule*.

scheduling mode. The type of scheduling operation set for the server and client. ADSM supports two scheduling modes for client operations: client-polling and server-prompted.

SCSI. Small computer system interface.

selective backup. A function that allows users to back up files or directories from a client domain that are not excluded in the include-exclude list and that meet the requirement for serialization as defined in the backup

copy group of the management class assigned to the files. Contrast with *incremental backup*.

sequential access media. Any volume that is accessed in a sequential manner, as opposed to a random manner. In ADSM, volumes are accessed sequentially if they reside in a device class other than DISK.

serialization. A copy group attribute that specifies what ADSM does if files are modified during backup or archive processing. The value of this attribute determines whether processing continues, is retried, or is stopped. See *static*, *dynamic*, *shared static*, and *shared dynamic*.

server. A program that provides services to other programs (clients).

server migration. The process of moving data from one storage pool to the next storage pool as controlled by the high and low migration thresholds. See *high migration threshold* and *low migration threshold*.

server options file. A file that specifies processing options for communication methods, tape handling, pool sizes, language, and date, time, and number formats.

server program. The program that provides backup, archive, space management, and administrative services to clients. The server program must be at the necessary level to provide all of these services.

server-prompted scheduling mode. A client/server communication technique where the server contacts the client when work needs to be done.

session resource usage. The amount of wait time, CPU time, and space used or retrieved during a client session.

shared dynamic. A copy group serialization value that specifies that a file must not be modified during a backup or archive operation. ADSM attempts to retry the backup or archive operation a number of times; if the file is in use during each attempt, ADSM will back up or archive the file on its last try even though the file is in use. See also *serialization*. Contrast with *dynamic*, *shared static*, and *static*.

shared static. A copy group serialization value that specifies that the file must not be modified during backup or archive. ADSM will retry the backup or archive operation a number of times; if the file is in use during each attempt, ADSM will not back up or archive the file. See also *serialization*. Contrast with *dynamic*, *shared dynamic*, and *static*.

shell. In the AIX and UNIX environments, a software interface between a user and the operating system of a computer. Shell programs interpret commands and

user interactions on devices such as keyboards, pointing devices, and touch-sensitive screens and communicate them to the operating system.

signal. (1) A simple method of communication between two processes. One process can inform the other process when an event occurs. (2) In operating system operations, a method of inter-process communication that simulates software interrupts.

signal handler. A subroutine called when a signal occurs.

SMIT. System Management Interface Tool.

SNA LU6.2. Systems Network Architecture Logical Unit 6.2.

socket. (1) An endpoint for communication between processes or applications. (2) A pair consisting of TCP port and IP address, or UDP port and IP address.

space management. The process of keeping sufficient free storage space available on a client node by migrating files to ADSM storage. The files are migrated based on criteria defined in management classes to which files are bound, and the include-exclude list. Synonymous with *hierarchical storage management*. See also *migration*.

space management client. A program that runs on workstations to allow users to maintain free space on their workstations by migrating and recalling files to and from ADSM storage. Synonymous with *hierarchical storage management client*.

SPACEMGPOOL. A disk storage pool defined by ADSM at installation. It can be the destination for files that are migrated from client nodes via ADSM space management. See *storage pool*.

stale copy status. Specifies that a volume copy is not available to the database or recovery log.

STANDARD copy group. A backup or archive copy group that is defined by ADSM at installation. See *copy group*.

STANDARD management class. A management class that is defined by ADSM at installation. See *management class*.

STANDARD policy domain. A policy domain that is defined by ADSM at installation. See *policy domain*.

STANDARD policy set. A policy set that is defined by ADSM at installation. See *policy set*.

stanza. A group of lines in a file that together have a common function or define a part of the system.

Stanzas are usually separated by blank lines or colons, and each stanza has a name.

startup window. A time period during which a schedule must be initiated.

static. A copy group serialization value that specifies that the file must not be modified during backup or archive. If the file is modified during the attempt, ADSM will not back up or archive the file. See also *serialization*. Contrast with *dynamic*, *shared dynamic*, and *shared static*.

storage hierarchy. A logical ordering of primary storage pools, as defined by an administrator with system privilege. Generally, the ordering is based on the speed and capacity of the devices that the storage pools use. In ADSM, the storage hierarchy is defined by identifying the *next* storage pool in a storage pool definition. See *storage pool*.

storage management services. A component that allows a central system to act as a file backup and archive server for local area network file servers and workstations.

storage pool. A named set of storage volumes that ADSM uses to store client data. A storage pool is either a primary storage pool or a copy storage pool. See *primary storage pool* and *copy storage pool*.

storage pool volume. A volume that has been assigned to an ADSM storage pool. See *volume*, *copy storage pool*, and *primary storage pool*.

storage privilege class. An administrative privilege class that allows an administrator to control the allocation and use of storage resources for the server, such as monitoring the database, recovery log, and data storage. Administrators can be authorized with unrestricted or restricted storage privilege. See *restricted storage privilege* or *unrestricted storage privilege*.

stub file. A file that replaces the original file on a client node when the file is migrated from the client node to ADSM storage.

superuser. See *root user*.

synchronized copy status. Specifies that the volume is the only volume copy or is synchronized with other volume copies in the database or recovery log. When synchronized, mirroring has started.

system privilege class. An administrative privilege class that allows an administrator to issue all server commands.

Systems Application Architecture (SAA). Software interfaces, conventions, and protocols that provide a

framework for designing and developing applications that are consistent across systems.

Systems Network Architecture (SNA). A set of rules for data to be transmitted in a network. Application programs communicate with each other using a layer of SNA called advanced program-to-program communications (APPC).

T

tape. A recording medium consisting of a long, narrow, flexible strip with a magnetic coating wound onto a reel or into a cartridge. See *cartridge* and *tape reel*.

tape library. (1) A term used to refer to a collection of tape cartridges. (2) An automated device that performs tape cartridge mounts and demounts without operator intervention.

Tape Library Dataserver. An automated tape library consisting of mechanical components, cartridge storage frames, IBM tape subsystems, and controlling hardware and software. The tape library dataserver performs tape cartridge mounts and demounts without operator intervention.

tape reel. A cylinder with flanges on which magnetic tape is wound. Devices such as the 3420 9-track tape device support tape reels. Contrast with *cartridge*.

tape volume prefix. A device class attribute that is the high-level-qualifier of the file name or the data set name in the standard tape label.

task help. A type of online help that provides a list of tasks that can be completed with a selected object. When you select a task, the help provides step-by-step information on how to complete the task.

TCP/IP. Transmission Control Protocol/Internet Protocol.

Telnet. In TCP/IP, the protocol that opens the connection to the system.

Transmission Control Protocol/Internet Protocol (TCP/IP). A set of communication protocols that support peer-to-peer connectivity functions for both local and wide area networks.

trusted communication agent. A program that performs communication tasks on behalf of the client or server, and ensures the security of the communications.

U

unit name. On an MVS server, a device class attribute that specifies a group of tape devices used with the MVS server. A unit name can be a generic device type, an esoteric unit name, or a physical device.

unrestricted policy privilege. An administrative privilege class that enables an administrator to manage policy objects for any policy domain.

unrestricted storage privilege. An administrative privilege class that enables an administrator to control the database, recovery log, and all storage pools.

utilization. The percent of assigned capacity used by the database or recovery log at a specific point of time.

V

validate. The process of ensuring that the active policy set contains a default management class and reporting on copy group definition errors.

version. The maximum number of backup copies retained for files and directories. The following copy group attributes define version criteria: versions data exists and versions data deleted.

Virtual Machine (VM). One of the family of IBM operating systems for the System/370 or System/390 processor, including VM/ESA, VM/XA, VM/SP, and VM/HPO. VM is one of the supported server environments.

Virtual Storage Extended (VSE). One of the family of IBM operating systems for the System/370 or System/390 processor, including VSE/ESA. VSE is one of the supported server environments.

VM. Virtual Machine.

volume. The basic unit of storage for the database, recovery log, or a storage pool. A volume can be an LVM logical volume, a standard file system file, a tape cartridge, or an optical cartridge. Each volume is identified by a unique volume identifier. See *database volume* and *storage pool volume*.

volume history file. A file that contains information about: volumes used for database backups and database dumps; volumes used for export of administrator, node, policy, or server data; and sequential access storage pool volumes that have been added, reused, or deleted. The information is a copy of the same types of volume information in the ADSM database.

volume set. An entire image of the database or recovery log, as displayed on the administrative graphical user interface.

VSE. Virtual Storage Extended.

W

WDSF/VM. Workstation Data Save Facility/Virtual Machine.

wildcard character. A character or set of characters used to specify an unknown number or set of characters in a search string. Also called *pattern-matching character*.

window. A part of a display screen with visible boundaries in which information is displayed.

windowed interface. A type of user interface that is either a graphical user interface or a text based inter-

face. The text based interface maintains a close affinity to the graphical user interface, including action bars and their associated pull-downs and windows. See *graphical user interface*.

workstation. A personal computer system capable of maintaining data files.

Workstation Data Save Facility/Virtual Machine (WDSF/VM). The predecessor product to ADSTAR Distributed Storage Manager.

X

X Windows. A network transparent windowing system developed by MIT. It is the basis for other products, such as Enhanced X Windows which runs on the AIX operating system.

Index

A

- accounting record, setting 407
- ACTIVATE POLICYSET command 38
- activating policy sets 38
- activity log
 - querying 252
 - retention period, setting 409
- administrative client
 - macro, using 17
 - mode
 - batch 13
 - console 12, 13
 - interactive 14
 - options
 - consolemode 15
 - id 15
 - itemcommit 15
 - noconfirm 15
 - outfile 15
 - password 15
 - quiet 15
 - session 11
- administrative command
 - components 3
 - entering 3
 - entry rules 4
 - help 224
 - issued by all administrators 33
 - name 3
 - naming conventions 8
 - parameters
 - keyword 3
 - positional 3
 - privilege class for 23, 217
 - processing
 - background 5
 - foreground 5
 - requiring analyst privilege 32
 - requiring operator privilege 31
 - requiring policy privilege 26
 - requiring system privilege 23
 - scheduling 424
 - using 35
 - wildcard characters 6
- administrative session, canceling 66
- administrator
 - adding 370
 - changing 380
 - deleting 377
 - exporting 195
 - importing 227

- administrator (*continued*)
 - locking 244
 - privilege class 23, 217
 - querying 256
 - registering 370
 - removing 377
 - renaming 380
 - unlocking 435
 - updating 437
- appending command output to a file 16
- application programming interface xii
- archive copy group
 - attributes 83
 - update 439
- archived file deletion, allowing 374, 464
- ASSIGN DEFMGMTCLASS command 40, 114
- assigning default management class 40, 114
- association, of client with schedule
 - defining 81
 - deleting 151
 - querying 260
- audit
 - client storage usage 43
 - data storage usage 43
 - database 544
 - license 43
 - occupancy 43
 - server storage usage 43
 - volume 45
- AUDIT LICENSES command 43
- authentication, setting 411
- authority
 - for command use 23
 - granting 217
 - revoking 401
 - unlocking 435

B

- BACKDELETE parameter 374
- background process
 - canceling 64
 - list of 64
 - querying 332
- backup
 - database 50
 - database volume 60
 - device information 54
 - file deletion, allowing 50
 - primary storage pool 56
 - scheduling for client 50
 - trigger for database 91

- backup (*continued*)
 - volume history information 60
- backup copy group
 - attributes 83
 - update 439
- BACKUP DB command 50
- BACKUP DEVCONFIG command 54
- BACKUP STGPOOL command 56
- BACKUP VOLHISTORY command 60
- batch mode
 - restrictions 13
 - starting and ending 13
 - using 13
- buffer pool statistics, resetting 386

C

- canceling
 - administrative session 66
 - background process 64
 - client node session 66
 - mount request 62
- changing
 - See updating
- client files, collocating 141, 144
- client node
 - adding 374
 - canceling session 66
 - changing 464
 - deleting 167
 - exporting 199
 - importing 230
 - locking 245
 - registering 374
 - removing 378
 - renaming 384
 - scheduling operations for 121, 424
 - session, canceling 66
 - unlocking 436
 - updating 464
- client schedule 121, 424
- collocation 141, 144
- command components 3
- command entry rules 4
- command line interface, using 11
- command macro 17
- command output, redirecting to a file 16
- command scheduling
 - defining 127
 - updating 482
- commands
 - ACTIVATE POLICYSET 38
 - ASSIGN DEFMGMTCLASS 40
 - AUDIT LICENSES 43
 - AUDIT VOLUME 45
 - BACKUP DB 50

- commands (*continued*)
 - BACKUP DEVCONFIG 54
 - BACKUP STGPOOL 56
 - BACKUP VOLHISTORY 60
 - CANCEL MOUNT 62
 - CANCEL PROCESS 64
 - CANCEL REQUEST 64
 - CANCEL SESSION 66
 - COMMIT 68
 - COPY DOMAIN 70
 - COPY MGMTCLASS 73
 - COPY POLICYSET 75
 - COPY SCHEDULE 77
 - DEFINE ASSOCIATION 81
 - DEFINE COPYGROUP 83
 - DEFINE DBBACKUPTRIGGER 91
 - DEFINE DBCOPY 95
 - DEFINE DBVOLUME 97
 - DEFINE DEVCLASS 99
 - DEFINE DOMAIN 107
 - DEFINE LOGCOPY 110
 - DEFINE LOGVOLUME 112
 - DEFINE MGMTCLASS 114
 - DEFINE POLICYSET 118
 - DEFINE SCHEDULE 121
 - DEFINE STGPOOL 135
 - DEFINE VOLUME 148
 - DELETE ASSOCIATION 151
 - DELETE COPYGROUP 153
 - DELETE DBBACKUPTRIGGER 156
 - DELETE DBVOLUME 158
 - DELETE DEVCLASS 160
 - DELETE DOMAIN 162
 - DELETE EVENT 164
 - DELETE FILESPACE 167
 - DELETE LOGVOLUME 170
 - DELETE MGMTCLASS 172
 - DELETE POLICYSET 174
 - DELETE SCHEDULE 176
 - DELETE STGPOOL 179
 - DELETE VOLHISTORY 181
 - DELETE VOLUME 185
 - DISABLE 188
 - DISMOUNT DEVICE 190
 - DISMOUNT VOLUME 191
 - ENABLE 192
 - EXPIRE INVENTORY 193
 - EXPORT ADMIN 195
 - EXPORT NODE 199
 - EXPORT POLICY 204
 - EXPORT SERVER 208
 - EXTEND DB 213
 - EXTEND LOG 215
 - GRANT AUTHORITY 217
 - HALT 221
 - HELP 224

commands (*continued*)

IMPORT ADMIN 227
IMPORT NODE 230
IMPORT POLICY 235
IMPORT SERVER 239
LOCK ADMIN 244
LOCK NODE 245
MACRO 246
MOVE DATA 249
QUERY ACTLOG 252
QUERY ADMIN 256
QUERY ASSOCIATION 260
QUERY AUDIT OCCUPANCY 262
QUERY CONTENT 265
QUERY COPYGROUP 270
QUERY DB 274
QUERY DBBACKUPTRIGGER 278
QUERY DBVOLUME 280
QUERY DEVCLASS 283
QUERY DOMAIN 287
QUERY EVENT 290
QUERY FILESPACE 300
QUERY LICENSE 304
QUERY LOG 307
QUERY LOGVOLUME 311
QUERY MGMTCLASS 314
QUERY MOUNT 318
QUERY NODE 320
QUERY OCCUPANCY 324
QUERY OPTION 327
QUERY POLICYSET 329
QUERY PROCESS 332
QUERY SCHEDULE 335
QUERY SESSION 339
QUERY STATUS 342
QUERY STGPOOL 347
QUERY VOLHISTORY 353
QUERY VOLUME 359
QUIT 365
REDUCE DB 366
REDUCE LOG 368
REGISTER ADMIN 370
REGISTER LICENSE 372
REGISTER NODE 374
REMOVE ADMIN 377
REMOVE NODE 378
RENAME ADMIN 380
RENAME FILESPACE 382
RENAME NODE 384
RESET BUFPOOL 386
RESET DBMAXUTILIZATION 388
RESET LOGCONSUMPTION 389
RESET LOGMAXUTILIZATION 390
RESTORE VOLUME 396
REVOKE AUTHORITY 401
ROLLBACK 405

commands (*continued*)

SET ACCOUNTING 407
SET ACTLOGRETENTION 409
SET AUTHENTICATION 411
SET EVENTRETENTION 412
SET LICENSEAUDITPERIOD 414
SET LOGMODE 416
SET MAXCMDRETRIES 419
SET MAXSCHEDESESSIONS 421
SET PASSEXP 423
SET QUERYSCHEDPERIOD 424
SET RANDOMIZE 426
SET REGISTRATION 428
SET RETRYPERIOD 430
SET SCHEDMODES 432
SET SERVERNAME 434
UNLOCK ADMIN 435
UNLOCK NODE 436
UPDATE ADMIN 437
UPDATE COPYGROUP 439
UPDATE DEVCLASS 450
UPDATE DOMAIN 458
UPDATE MGMTCLASS 460
UPDATE NODE 464
UPDATE POLICYSET 467
UPDATE SCHEDULE 469
UPDATE STGPOOL 482
UPDATE VOLHISTORY 493
UPDATE VOLUME 495
VALIDATE POLICYSET 501
VARY 504

commands in a macro
 committing 20
 rolling back 20

commands issued by
 all administrators 33
 analyst administrators 32
 operator administrators 31
 policy administrators 27
 storage administrators 29
 system administrators 23

COMMIT command 15, 18, 20
committing commands in a macro 20
console mode, using 12
continuation characters, using 16
COPY DOMAIN command 70
copy group
 defining 83
 deleting 153
 querying 270
 updating 439
COPY MGMTCLASS command 73
COPY POLICYSET command 75
COPY SCHEDULE command 77
copying
 domain 70

copying (*continued*)
 management class 40
 policy domain 70
 policy set 75
 schedule 77

D

damaged file, restore 391
damaged volume, restore 396, 515
data, moving 249
database
 auditing 544
 backup 50
 backup trigger 91
 backup volume information 280, 508
 dumping 530
 extending 213
 installing 536
 querying 274
 reducing 366
 reloading 538
 restoring 416
 verifying 544
 volume information 280, 511
database backup series 50, 91
database maximum utilization statistic, resetting 388
database recovery
 DSMSERV RESTORE DB 515
 log information 515
 log, mode for 515
 point-in-time 515
 roll-forward 515
database salvage utility commands
 DSMSERV AUDITDB 544
 DSMSERV DUMPDB 530
 DSMSERV INSTALL 536
 DSMSERV LOADDDB 538
database volume
 defining 97
 deleting 158
 querying 280, 511
default management class, assigning 40
DEFINE ASSOCIATION command 81
DEFINE COPYGROUP command 83
DEFINE DBBACKUPTRIGGER command 91
DEFINE DBCOPY command 95
DEFINE DBVOLUME command 97
DEFINE DEVCLASS command 99
DEFINE DOMAIN command 107
DEFINE LOGCOPY command 110
DEFINE LOGVOLUME command 112
DEFINE MGMTCLASS command 114
DEFINE POLICYSET command 118
DEFINE SCHEDULE command 121

DEFINE STGPOOL command 135
DEFINE VOLUME command 148
defining
 association 81
 copy group 83
 database backup trigger 91
 database volume 97
 device class 99
 domain 107
 management class 114
 mirrored database volume 95
 mirrored log copy 110, 118
 policy set 118
 recovery log volume 112
 schedule 121
 storage pool 135
 storage pool volume 148
DELETE ASSOCIATION command 151
DELETE COPYGROUP command 153
DELETE DBBACKUPTRIGGER command 156
DELETE DBVOLUME command 158
DELETE DEVCLASS command 160
DELETE DOMAIN command 162
DELETE EVENT command 164
DELETE FILESPACE command 167
DELETE LOGVOLUME command 170
DELETE MGMTCLASS command 172
DELETE POLICYSET command 174
DELETE SCHEDULE command 176
DELETE STGPOOL command 179
DELETE VOLHISTORY command 181
DELETE VOLUME command 185
deleting
 archived files 374
 backed-up files 374
 copy group 153
 database backup trigger 156
 database volume 158
 device class 160
 domain 162
 event record 164
 file space 167
 management class 172
 policy set 174
 recovery log volume 170
 schedule 176
 storage pool 179
 storage pool volume 185
 volume history 181
destroyed volume, restore 396, 515
device class
 defining 99
 deleting 160
 querying 283
 updating 450

- device configuration file 54
- DISABLE command 188
- disabling the server 188
- DISMOUNT VOLUME command 191
- dismounting volumes 191
- documentation, user xi
- domain
 - copying 70
 - defining 107
 - deleting 162
 - querying 287
 - updating 458
- DSMSERV AUDITDB command 529, 544
- DSMSERV DISPLAY DBBACKUPVOLUME command 507, 508
- DSMSERV DISPLAY DBVOLUMES command 507, 511
- DSMSERV DISPLAY LOGVOLUMES command 507, 513
- DSMSERV DUMPDB command 507, 529, 530
- DSMSERV INSTALL command 529, 536
- DSMSERV LOADDB command 529, 538
- DSMSERV RESTORE DB command 507, 515
- dumping the database 530
 - See also* backup

E

- ENABLE command 192
- enabling the server 192
- ending
 - batch mode 13
 - console mode 12
 - interactive mode 15
- entering administrative commands 3
- event record, deleting 164
- event, querying 290
- expiration
 - date, for volume 85
 - inventory 193
 - password, interval 423
- EXPIRE INVENTORY command 193
- expiring inventory 193
- EXPORT ADMIN command 195
- EXPORT NODE command 199
- EXPORT POLICY command 204
- EXPORT SERVER command 208
- export, administrator 195
- EXTEND DB command 213
- EXTEND LOG command 215
- extending
 - database 213
 - recovery log 215

F

- file space
 - deleting 167
 - querying 252
 - renaming 382
 - storage used by 324
- files
 - collocating for client node 141
 - deleting expired files 193
 - moving 249
- formatting command output on a terminal 6

G

- GRANT AUTHORITY command 217
- granting, authority 217

H

- HALT command 221
- halting the server 221
- HELP command 224
- help for administrative commands 3
- hierarchical storage management 114, 460
- hierarchy, storage pool
 - considerations for backup 56
 - defining 135, 484
- HSM
 - See* hierarchical storage management

I

- IMPORT ADMIN command 227
- IMPORT NODE command 230
- IMPORT POLICY command 235
- IMPORT SERVER command 239
- importing
 - administrator 227
 - node 230
 - policy information 235
 - server 239
- incremental backup of database 50, 91
- installing
 - database 536
 - recovery log 536
- interactive mode
 - continuation characters 16
 - ending 16
 - quitting 365
 - restrictions 14
 - starting 14
 - using 14
- inventory, expiring 193

K

keyword parameters 3

L

license

- auditing 43
- querying 304
- registering 372

LICENSEAUDITPERIOD, setting 414

LOCK ADMIN command 244

LOCK NODE command 245

locking

- administrator 244
- node 245

log mode, recovery 416

M

macro

- commit individual commands 20
- continuation characters 16
- controlling command processing 20
- running commands 20
- scheduling for client 122
- substitution variables 19
- testing 20
- using 17
- writing 18

MACRO command 246

management class

- copying 73
- defining a default 40
- deleting 172
- querying 314
- updating 460

maximum retries for client schedule, setting 419

migration

- files from client node 114, 460
- storage pool 56

mirrored database volume, defining 95

mirrored log volume, defining 112

monitoring, ADSTAR Distributed Storage Manager activities 12

mount request, canceling 64

mounted sequential access volumes

- dismounting 191
- querying 318

MOVE DATA command 249

moving files 249

N

name

- administrator 380
- case-sensitivity 8

name (*continued*)

- client 8, 384
- file space 8, 382
- server 8

naming

- conventions 8
- restrictions 8
- volumes 8

node

- associating with a schedule 424
- changing 464
- data storage usage 262
- deleting 378
- exporting 199
- importing 230
- locking 245
- querying 320
- registering 374
- removing 378
- renaming 384
- unlocking 436
- updating 464

normal, recovery log mode 50

O

occupancy, querying 324

offsite volume 149, 495, 497

options, server

- importance in recovery 50
- querying 327

P

parameters

- keyword 3
- positional 3

password

- authentication, setting 411
- expiration interval 423
- UPDATE ADMIN command 437

pattern-matching expression, creating 6

policy domain

- copying 70
- defining 107
- deleting 162
- querying 287
- updating 458

policy set

- activating 38
- copying 75
- defining 118
- deleting 174
- querying 329
- updating 467
- validating 501

- polling, SET QUERYSCHEDPERIOD 424
- printing, redirected output 16
- privilege class
 - administrative commands 29
 - administrator issued 33
 - all administrators 33
 - analyst 32
 - assigning 217
 - none 33
 - operator 23, 31
 - policy 27
 - restricted 29
 - storage 23, 29
 - system 23
 - system administrator 23
 - unrestricted 29
 - using commands 23
- process
 - background 64
 - canceling 64
 - querying 332
- program directory xi
- publications xi

Q

- QUERY ACTLOG command 252
- QUERY ADMIN command 256
- QUERY ASSOCIATION command 260
- QUERY AUDITOCUPANCY command 262
- QUERY CONTENT command 265
- QUERY COPYGROUP command 270
- QUERY DB command 274
- QUERY DBBACKUPTRIGGER command 278
- QUERY DBVOLUME command 280
- QUERY DEVCLASS command 283
- QUERY DOMAIN command 287
- QUERY LICENSE command 290, 304
- QUERY LOG command 307
- QUERY LOGVOLUME command 311
- QUERY MGMTCLASS command 314
- QUERY MOUNT command 318
- QUERY NODE command 320
- QUERY OCCUPANCY command 324
- QUERY OPTION command 327
- QUERY POLICYSET command 329
- QUERY PROCESS command 332
- QUERY SCHEDULE command 335
- QUERY SESSION command 339
- QUERY STATUS command 342
- QUERY STGPOOL command 347
- QUERY VOLHISTORY command 353
- QUERY VOLUME command 353
- querying
 - activity log 252
 - administrator 256

- querying (*continued*)
 - audit occupancy 262
 - background processing 5
 - clients with schedules 260
 - completed events 290
 - contents of a volume 265
 - copy group 270
 - database 274
 - database backup volume 508
 - database volume 280, 511
 - device class 283
 - domain 287
 - file space 300
 - license 304
 - management class 314
 - node 320
 - occupancy 324
 - option 327
 - policy set 329
 - process 332
 - recovery log 307
 - recovery log volume 311, 513
 - scheduled events 290
 - schedules 335
 - session 339
 - status 342
 - storage pool 347
 - storage pool volume 359
 - volume history file 353
- QUIT command 365
- quitting 365

R

- reclaim, volumes 144
- recovery log
 - backup trigger 91
 - consumption statistic, reset 389
 - extending 215
 - installing 536
 - log volume
 - defining 112
 - deleting 170
 - displaying 513
 - querying 311
 - maximum utilization statistic, resetting 390
 - mode for saving records
 - NORMAL 416
 - ROLLFORWARD 416
 - reducing 368
- redirecting
 - appending output 16
 - characters 16
 - command output 16
 - writing output 16

- REDUCE DB command 366
- REDUCE LOG command 368
- reducing
 - database 366
 - recovery log 368
- REGISTER ADMIN command 370
- REGISTER LICENSE command 372
- REGISTER NODE command 374
- registering
 - administrator 370
 - license 372
 - node 374
- registration, setting 428
- reloading the database 538
- REMOVE ADMIN command 377
- removing
 - administrator 377
 - client association 151
 - node 378
- RENAME ADMIN command 380
- RENAME FILESPACE command 382
- RENAME NODE command 384
- renaming
 - administrator 380
 - filesystem 382
 - node 384
- RESET BUFPOOL command 386
- RESET DBMAXUTILIZATION command 388
- RESET LOGCONSUMPTION command 389
- RESET LOGMAXUTILIZATION command 390
- resetting statistic
 - buffer pool 386
 - maximum utilization (database) 388
 - maximum utilization (recovery log) 390
 - recovery log consumption 389
- RESTORE STGPOOL command 391
- RESTORE VOLUME command 396
- retention period, setting 412
- retry period, setting 430
- REVOKE AUTHORITY command 401
- revoking authority for administrator 401
- roll-forward recovery
 - description 50
 - of database 515
 - of database volume 516
- ROLLBACK command 405
- rolling back, commands in a macro 11

S

schedule

- administrative command 121, 469
- changing 469
- client node, associate with 81
- client operation 121, 424
- copying 77

schedule (*continued*)

- defining 114
- deleting 176
- parameters
 - administrative 128
 - client 122
 - querying 335
 - querying results of 290
 - retries, setting number 127
 - retries, time between 127
 - updating 469
- scheduling mode, setting 432
- scratch volumes in storage pool 142, 145
- server
 - disabling 188
 - enabling 192
 - exporting 208
 - halting 221
 - importing 239
- server console
 - restrictions 11
 - using 11
- server name, setting 434
- session
 - maximum number scheduled 421
 - querying client 339
- SET ACCOUNTING command 407
- SET ACTLOGRETENTION command 409
- SET AUTHENTICATION command 411
- SET EVENTRETENTION command 412
- SET LICENSEAUDITPERIOD command 414
- SET LOGMODE command 416
- SET MAXCMDRETRIES command 419
- SET MAXSCHEDESESSIONS command 421
- SET PASSEXP command 423
- SET QUERYSCHEDPERIOD command 424
- SET RANDOMIZE command 426
- SET REGISTRATION command 428
- SET RETRYPERIOD command 430
- SET SCHEDMODES command 432
- SET SERVERNAME command 434
- setting
 - accounting record 407
 - authentication 411
 - frequency for client-polling 424
 - license audit period 414
 - maximum retries of a command 419
 - maximum scheduled sessions 421
 - password 423
 - randomization of start times 426
 - recovery log mode 416
 - registration 428
 - retention period 409
 - retention period for event records 412
 - retry period 430
 - scheduling mode 432

- setting (*continued*)
 - server name 434
- space management
 - DEFINE MGMTCLASS command 114
 - UPDATE MGMTCLASS command 460
- starting
 - batch mode 13
 - using console mode 12
 - using interactive mode 14
- status, querying 342
- storage management services xi
- storage pool
 - defining 135
 - deleting 179
 - querying 347
 - updating 482
- storage pool volume
 - auditing 45
 - defining 148
 - deleting 185
 - querying 265, 359
 - random access 135, 136
 - scratch 142, 145
 - sequential access 139
 - updating 495
 - varying 504
- storage volumes, naming 3
- syntax diagrams, using 1
- system privilege class, commands 23

T

- thresholds, storage pool migration 136
- trademarks x
- trigger, database backup 91

U

- UNLOCK ADMIN command 435
- UNLOCK NODE command 436
- unlocking
 - administrator 435
 - node 436
- UPDATE ADMIN command 437
- UPDATE COPYGROUP command 439
- UPDATE DBBACKUPTRIGGER command 447
- UPDATE DEVCLASS command 450
- UPDATE DOMAIN command 458
- UPDATE MGMTCLASS command 460
- UPDATE NODE command 464
- UPDATE POLICYSET command 467
- UPDATE SCHEDULE command 469
- UPDATE STGPOOL command 482
- UPDATE VOLUME command 495
- updating
 - administrator 437

- updating (*continued*)
 - copy group 439
 - device class 450
 - domain 458
 - management class 460
 - node 464
 - policy set 467
 - schedule 469
 - storage pool 482
 - storage pool volume 495
- user documentation xi
- using
 - administrative client options 15
 - administrative privilege classes 23
 - command line interface 11
 - commands by administrative privilege 23
 - continuation characters 16
 - macros 17
 - redirection 16
 - redirection characters 16
 - syntax diagram 1

V

- VALIDATE POLICYSET command 501
- validating, policy set 501
- VARY command 504
- varying, volume 504
- volume history
 - backing up 60
 - deleting 181
 - querying 353
- volume, storage
 - defining 148
 - naming 3
 - random access 148
 - reclaiming 141
 - scratch
 - maximum in storage pool 142, 145
 - using 148
 - sequential access 148

W

- wildcard, match-any characters
 - match-any characters 7
 - match-exactly-one 7
 - using 6
- writing, command output to a file 16

Communicating Your Comments to IBM

ADSTAR Distributed Storage Manager
for MVS
Administrator's Reference
Version 2

Publication No. SH26-4040-00

If you especially like or dislike anything about this book, please use one of the methods listed below to send your comments to IBM. Whichever method you choose, make sure you send your name, address, and telephone number if you would like a reply.

Feel free to comment on specific errors or omissions, accuracy, organization, subject matter, or completeness of this book. However, the comments you send should pertain to only the information in this manual and the way in which the information is presented. To request additional publications, or to ask questions or make comments about the functions of IBM products or systems, you should talk to your IBM representative or to your IBM authorized remarketer.

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

If you are mailing a readers' comment form (RCF) from a country other than the United States, you can give the RCF to the local IBM branch office or IBM representative for postage-paid mailing.

- If you prefer to send comments by mail, use the RCF at the back of this book.
- If you prefer to send comments by FAX, use this number:
 - United States and Canada: 520 799-2906
 - Other countries: (1) 520 799-2906

The contact department is 61C/031.

- If you prefer to send comments by electronic mail, use one of the following addresses:
 - Internet: starpubs@vnet.ibm.com (or starpubs at vnet.ibm.com)
 - IBMLink from U.S.A.: STARPUBS at SJEVM5
 - IBMLink from Canada: STARPUBS at TORIBM
 - IBM Mail Exchange: USIB3VVD at IBMMAIL

Make sure to include the following in your note:

- Title and publication number of this book
- Page number or topic to which your comment applies.

Readers' Comments — We'd Like to Hear from You

**ADSTAR Distributed Storage Manager
for MVS
Administrator's Reference
Version 2**

Publication No. SH26-4040-00

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

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

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

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

Please tell us how we can improve this book:

Thank you for your responses. May we contact you? Yes No

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

Name

Address

Company or Organization

Phone No.




Cut or Fold
Along Line


Fold and Tape

Please do not staple

Fold and Tape



NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES



BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

International Business Machines Corporation
Information Development
Department 61C
9000 South Rita Road
TUCSON AZ 85775-4401



Fold and Tape

Please do not staple

Fold and Tape

Cut or Fold
Along Line



Program Number: 5655-119



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

SH26-4040-00



Spine information:



**ADSTAR Distributed Storage Manager
for MVS**

Administrator's Reference

Version 2