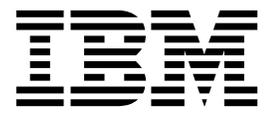


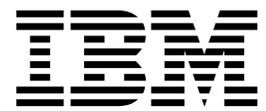
IBM Spectrum Protect for Space Management
Version 8.1.2

User's Guide



IBM Spectrum Protect for Space Management
Version 8.1.2

User's Guide



Note:

Before you use this information and the product it supports, read the information in “Notices” on page 193.

This edition applies to version 8, release 1, modification 2 of IBM Spectrum Protect for Space Management (product number 5725-X04), and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright IBM Corporation 1995, 2017.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Figures vii

Tables ix

About this publication xi

Who should read this publication xi

Publications xi

Conventions used in this publication xi

**New for IBM Spectrum Protect for
Space Management xiii**

**Chapter 1. Overview of the space
management client 1**

Migrating files overview 2

Recalling migrated files overview 3

Reconciling file systems overview 4

**Chapter 2. Installing the space
management client 5**

General installation prerequisites and considerations 5

Installation overview for the space management
client for AIX GPFS systems 7

Environment requirements for the space
management client for AIX GPFS systems 9

Installing the space management client for AIX
systems 9

Upgrading the space management client for AIX
GPFS systems 10

Uninstalling the space management client for AIX
GPFS systems. 11

Installation overview for the space management
client for Linux GPFS systems 12

Environment requirements for the space
management client for Linux GPFS systems . . 14

Installing the space management client for Linux
GPFS systems 14

Upgrading the space management client for
Linux GPFS systems 15

Uninstalling the space management client for
Linux GPFS systems 16

Registering the space management client with the
IBM Spectrum Protect server 17

Registering your node by using open registration 17

Closed registration 18

Password management 18

**Chapter 3. Configuring the space
management client 19**

Displaying space management client options . . . 20

Editing the options files by using the
backup-archive client 20

Editing the options file dsm.sys. 20

Editing the options file dsm.opt 22

Configuring the space management client to connect
to a secondary IBM Spectrum Protect server . . . 23

Optional setup features 24

Assignment of management classes to files . . . 24

Displaying management class information . . . 26

Include-exclude file options 26

Environment variables. 28

LAN-free data transfer for HSM 29

Configuring HSM on GPFS clusters 31

Logs for HSM activity and error messages . . . 32

**Chapter 4. Adding and configuring
space management for a file system . . 35**

Adding space management to file systems 36

Adding space management to nested file systems . 36

Adding space management to an exported file
system 37

Removing or reconfiguring a space-managing node
in a GPFS cluster 37

Changing the mount point of a space-managed file
system 38

Mounting a space-managed file system at a
mount point for which a file space does not exist. 38

Mounting a previously space-managed file
system at a new mount point for which a file
space already exists. 38

Adding space management to workload partitions
on AIX V6.1 and V7.1 operating systems 39

Configuration by using the command line 40

Settings that control the space usage of a file system 40

Minimum migration file size 42

Migration threshold percentages 42

Premigration percentage 43

Quotas 43

Stub file size 44

Minimum stream file size. 44

Minimum partial file size. 45

Maximum number of files 45

Managing a file system with multiple IBM Spectrum
Protect servers 45

Enabling a file system to be managed by multiple
IBM Spectrum Protect servers 46

Adding space management to a file system in a
multiple-server environment. 47

Removing IBM Spectrum Protect servers from a
multiple-server environment. 48

Limitations for multiple-server environments . . 49

Replicating nodes in a multiple-server
environment 50

Deactivating space management 51

Reactivating space management 51

Removing space management 52

Chapter 5. Migrating files 53

Migration types	53
File migration eligibility	54
File premigration	55
Automatic file migration	55
Candidate selection for automatic migration	56
Manually starting threshold migration	57
Selective file migration	58
Migrating selectively by using the dsmmigrate command	58
Migration of a single file system to two or more IBM Spectrum Protect servers	59

Chapter 6. Backup and restore on space managed file systems 61

Archive and retrieve of migrated files by using the backup-archive client	62
Migration status after a migrated file is archived	62
Migration status after a migrated file is archived and retrieved	63
Back up before migration	63
Backing up migrated files	63
Restoring migrated files	65
Re-create stub files by using the dsmmigundelete command	66
Restoring space managed file systems	67
Restoring a file system that is backed up and migrated to the same IBM Spectrum Protect server	68
Restoring a file system after backup and migration operations used different IBM Spectrum Protect servers	68
Restoring a file system in an environment that is managed by multiple IBM Spectrum Protect servers	69

Chapter 7. Recalling migrated files. 71

Transparent recall processing	71
Selective recall processing	72
Normal recall mode	72
Partial file recall mode.	72
Streaming recall mode.	73
How HSM determines which recall mode to use	74
Setting the recall mode by using the dsmattr command	74
Optimized tape recall processing	75
List files for optimized tape recalls	76
Recalling files with optimization for tape	78

Chapter 8. File system reconciliation 79

Automatic reconciliation	79
Manual reconciliation of file systems	80
Orphaned stub files	81
Configuring immediate reconciliation.	81
Reconciling by using a GPFS policy	82

Chapter 9. Space management daemons 83

The space monitor daemon	83
------------------------------------	----

The recall daemon	84
The scout daemon	84
The watch daemon	85
Stopping the space management daemons	85

Chapter 10. Scheduled HSM service . . . 87

Chapter 11. Options files reference . . . 89

afmskipuncachedfiles	89
candidatesinterval	90
checkfororphans	90
checkthresholds	91
compression	92
defaultserver	92
errorlogname	93
errorprog	94
hsmdisableautomigdaemons	94
hsmdistributedrecall	95
hsmenableimmediatemigrate	95
hsmeventdestroy	96
hsmextobjidattr	97
hsmgroupedmigrate	97
hsmlogeventflags	98
hsmlogmax	99
hsmlogname	99
hsmlogretention	100
hsmlogsampleinterval	101
hsmmaxrecalltapedrives	101
hsmmigeroblockfiles	102
hsmmultiserver	103
includexcl	103
maxcandprocs	104
maxmigrators	104
maxrecalldaemons	105
maxthresholdproc	105
migfileexpiration	106
migrateserver	106
minmigfilesize	107
minrecalldaemons	108
reconcileinterval	108
restoremigstate	109
skipmigrated	110

Chapter 12. HSM client command reference 113

Standard option formats	113
Help for commands	114
Displaying file and file system information	114
Client return codes	115
HSM command summary	116
dmkilld	118
dsmattr	118
dsmautomig	121
dsmdf	123
dsmdu	124
dsmls	126
dsmmigfs add and update	128
dsmmigfs addmultiserver, querymultiserver, and removemultiserver	134
dsmmigfs control	135

dsmmigfs deactivate, reactivate, and remove	137
dsmmigfs globaldeactivate and globalreactivate	138
dsmmigfs help	139
dsmmigfs query	139
dsmmigfs rollback	142
dsmmigfs enablefailover, and disablefailover	143
dsmmigfs stop, start, and restart	144
dsmmigfs takeover	146
dsmmighelp	147
dsmmigquery	147
dsmmigrate	149
dsmmigundelete	153
dsmmonitord	155
dsmq	156
dsmrecall	157
dsmrecalld	160
dsmreconcile	161
dsmrm	163
dsmscoutd	164
dsmsetpw	165
dsmwatchd	166

Chapter 13. HSM GPFS clients Perl

scripts	169
dsmMultiServerUpgrade.pl	169
dsmNextServername.pl	169
dsmreconcileGPFS.pl	170
dsmRemoveServer.pl	171

Chapter 14. Troubleshooting the space management client

Checking the time stamp of a daemon process-ID file	173
---	-----

Creating a dump file	174
HSM entries in the syslog file	174
Client trace classes for the space management client	176
Enabling a Space Management client trace	177
Enabling daemon traces	178
Tracing the watch daemon	179
Tracing the recall daemon	179
Tracing the scout daemon	180
Tracing the monitor daemon	181
Log files to aid troubleshooting	181
Commands for displaying HSM information	181
Commands for displaying GPFS information	183
Commands for displaying operating system information	183
Other tools for collecting information	184
HSM common problems and solutions	184

Appendix A. Control files in the .SpaceMan directory	187
Space requirement for HSM control files	188

Appendix B. Accessibility features for the IBM Spectrum Protect product family	191
---	-----

Notices	193
----------------	-----

Glossary	197
-----------------	-----

Index	199
--------------	-----

Figures

1. Overview of hierarchical storage management	1	3. Generation of list files	78
2. Ordered list file generation	76		

Tables

1. The space management client for AIX GPFS systems installation packages	7	13. space management client determining recall modes	74
2. Linux x86_64 GPFS available packages	12	14. Automatic reconciliation tasks	79
3. Linux z Systems GPFS available packages	12	15. Option format examples: options without values	113
4. Package names and descriptions	15	16. Option format examples: options with values	114
5. space management client options in the dsm.opt file	22	17. HSM commands to display file and file system information	114
6. The space management client attributes in a management class	25	18. An explanation of client return codes	115
7. Assigning management classes to your files	25	19. HSM command summary	116
8. Include and Exclude Statements.	26	20. Space management for locally and remotely mounted file systems	125
9. Environment Variables	28	21. Space management for locally and remotely mounted file systems	126
10. LAN-free data transfer options	30	22. Trace classes	176
11. dsmmigundelete command actions without the expiring option	67	23. Common HSM problems and resolutions	184
12. dsmmigundelete command actions with the expiring option	67	24. Control files stored in the .SpaceMan directory	187

About this publication

This publication provides the information to install, configure, monitor, and troubleshoot problems with IBM Spectrum Protect™ for Space Management.

Who should read this publication

This publication is intended for those who are responsible for installing, setting up, and administering the IBM Spectrum Protect for Space Management client. In this publication, it is assumed that you have a working knowledge of IBM Spectrum Protect for Space Management.

Publications

The IBM Spectrum Protect product family includes IBM Spectrum Protect Snapshot, IBM Spectrum Protect for Space Management, IBM Spectrum Protect for Databases, and several other storage management products from IBM®.

To view IBM product documentation, see IBM Knowledge Center.

Conventions used in this publication

This publication uses the following typographical conventions:

Example	Description
autoexec.ncf hsmgui.exe	A series of lowercase letters with an extension indicates program file names.
DSMI_DIR	A series of uppercase letters indicates return codes and other values.
dsmQuerySessInfo	Boldface type indicates a command that you type on a command line, the name of a function call, the name of a structure, a field within a structure, or a parameter.
<i>timeformat</i>	Boldface italic type indicates a backup-archive client option. The bold type is used to introduce the option, or used in an example.
<i>dateformat</i>	Italic type indicates an option, the value of an option, a new term, a placeholder for information you provide, or for special emphasis in the text.
maxcmdretries	Monospace type indicates fragments of a program or information as it might appear on a display screen, such a command example.
plus sign (+)	A plus sign between two keys indicates that you press both keys at the same time.

New for IBM Spectrum Protect for Space Management

Learn about new features and updates for IBM Spectrum Protect for Space Management Version 8.1.2.

For a list of new features and updates in this release and previous Version 8.1 releases, see IBM Spectrum Protect for Space Management updates.

New and changed information is indicated by a vertical bar (|) in the margin.

Chapter 1. Overview of the space management client

The IBM Spectrum Protect for Space Management client migrates files from your local file system to IBM Spectrum Protect server storage and can then recall the files either automatically or selectively. Migrating files to storage frees space for new data on your local file system and takes advantage of lower-cost storage resources that are available in your network environment.

The figure shows an overview of hierarchical storage management.

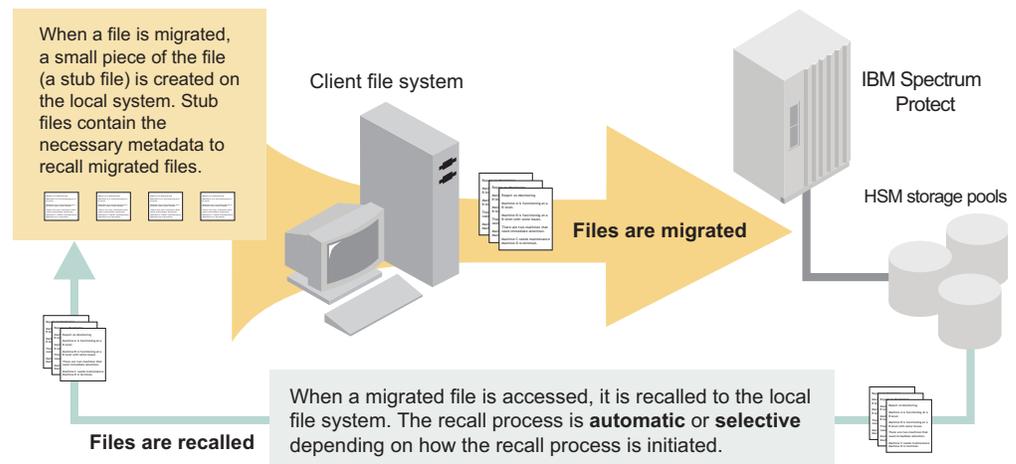


Figure 1. Overview of hierarchical storage management

When a file is migrated from your local system to IBM Spectrum Protect server storage, a placeholder, or stub file, is created in place of the original file. Stub files contain the necessary information to recall your migrated files and remain on your local file system. The migrated files appear to be on your local file system. This process contrasts with archiving, where you usually delete files from your local file system after you archive them.

The space management client provides space management services for locally mounted file systems, and it migrates regular files only. It does not migrate character special files, block special files, named pipe files, or directories.

File migration, unlike file backup, does not protect against accidental file deletion, file corruption, or disk failure. Continue to back up your files regardless of whether they are on your local file system or are migrated to IBM Spectrum Protect server storage. The IBM Spectrum Protect backup-archive client backs up and restores resident and migrated files. If you accidentally delete stub files from your local file system, or if you lose your local file system, you can restore the stub files or the complete files.

For planned processes, such as storing a large group of files in storage and returning them to your local file system for processing, use the archive and retrieve processes. The backup-archive client is used to archive and retrieve copies of migrated files.

The space management client functions for threshold migration, demand migration, selective migration, selective and transparent recall includes processing GPFS™ file systems that contain multiple space-managed storage pools.

The space management client has commands that you can run from a shell. You can also use the commands in scripts and cron jobs.

For example, the following commands migrate all files that are owned by user ibm:

```
find /hsmmanagedfilesystem -user ibm -print > /tmp/filelist  
dsmmigrate -filelist=/tmp/filelist
```

Your IBM Spectrum Protect server administrator associates management classes with your files. You, as a root user, can do the following tasks:

- Select space management options and settings.
- Assign management classes to your files.
- Exclude files from space management.
- Schedule space management services.

You can set the following options for migration:

- Which files are eligible for automatic migration
- The order in which files are migrated
- Where the migrated files are stored
- How much free space is maintained on your local file system

You prioritize files for migration by their file size or by the number of days since your files were last accessed.

Related concepts:

Chapter 2, “Installing the space management client,” on page 5

“File migration eligibility” on page 54

Chapter 6, “Backup and restore on space managed file systems,” on page 61

Chapter 7, “Recalling migrated files,” on page 71

Chapter 8, “File system reconciliation,” on page 79

Migrating files overview

The space management client provides both automatic and selective migration. After file migration begins, the space management client sends a copy of your file to IBM Spectrum Protect server storage volumes on disk devices or devices that support removable media, such as tape and replaces the original file with a stub file on your local file system.

The stub file is a small replacement file that makes it appear as though the original file is on the local file system. It contains required information to locate and recall a migrated file and to respond to specific UNIX commands without recalling the file.

Automatic migration monitors space usage and automatically migrates eligible files according to the options and settings that you select. The space management client provides two types of automatic migration: threshold migration and demand migration.

Threshold migration maintains a specific level of free space on your local file system. When space usage reaches the high threshold that you set for your file system, eligible files are migrated to IBM Spectrum Protect server storage automatically. When space usage drops to the low threshold that you set for your file system, file migration stops.

Demand migration responds to an out-of-space condition on your local file system. Demand migration starts automatically if your file system runs out of space. On GPFS file systems, the process receives an out-of-space (ENOSPC) return code and stops.

On GPFS file systems, you can use the GPFS policy engine to monitor space thresholds and search for migration candidates.

For information about configuring GPFS integration with the space management client, see Technote 7018848.

Selective migration moves specific files from your local file system to IBM Spectrum Protect server storage. For example, if you know that you will not be using a particular group of files for an extended time, you can migrate them to IBM Spectrum Protect server storage to free additional space on your local file system.

Related concepts:

“File migration eligibility” on page 54

Recalling migrated files overview

You can recall a migrated file to your local file system from IBM Spectrum Protect server storage either selectively or transparently. Files are recalled in either normal, partial, or streaming mode.

Selective recall returns specified, migrated files to your local file system. You select the files that you want to recall. When you selectively recall a file, you store it in its originating file system. Selective recall overrides the recall mode that you set for a migrated file with normal recall mode.

Transparent recall automatically returns a migrated file to your local file system when you access the file. If you change the recall mode for a migrated file, you change how the space management client recalls a migrated file.

Normal recall mode recalls a migrated file to its originating file system. The recalled file remains on your local file system. When you close the unmodified file, the copy that currently resides in IBM Spectrum Protect server storage remains valid. The local copy is premigrated.

Optimized tape recall optimizes the tape access during a recall of a list of files. Files indicated with the **filelist** option to the **dsmrecall** command are recalled in an order that optimizes tape access. The recall order minimizes tape mount and unmount operations and minimizes tape seek operations.

Note: The following recall modes apply only to read operations. For write and truncate operations on migrated files, the normal recall mode is always used.

Partial file recall mode recalls a portion of a migrated file and is valid only for GPFS systems. A partial recall avoids recalling an entire file, when only a small portion of the file is required by an application. The space management client intercepts a read request for a file configured for partial file recall. The space management client then calculates which portion of the file to recall based on the offsets contained in the read request. This results in time and disk space savings, since only a portion of the file is recalled.

Streaming recall mode enables or disables an asynchronous recall of migrated files. The recalled portion of the file can be accessed while the file is recalled. Streaming recall mode is valid for read-only operations on the file.

Note: Partial file recall mode takes precedence over streaming recall mode.

Related concepts:

Chapter 7, “Recalling migrated files,” on page 71

“Selective recall processing” on page 72

“Optimized tape recall processing” on page 75

“Partial file recall mode” on page 72

“Streaming recall mode” on page 73

Reconciling file systems overview

When you modify or delete a migrated or premigrated file on the local file system, the migrated copy in IBM Spectrum Protect server storage becomes obsolete. During reconciliation, any obsolete copies of migrated or premigrated files are marked for expiration. When the copies expire, they are removed from the server.

The default interval for reconciliation is 24 hours. When reconciliation is not driven by GPFS policy, the root user can set the reconcile interval with the **reconcileinterval** option. If you have many space-managed file systems on a system, increase this value to reduce the impact that the **dsmreconcile** command might have on system performance.

Related concepts:

Chapter 8, “File system reconciliation,” on page 79

Chapter 2. Installing the space management client

You install the space management client on your workstation and register it as a client node with the IBM Spectrum Protect server. The files on your node must be associated with a management class on the IBM Spectrum Protect server that is configured for space management. Read both the general and system-specific requirements before you install the space management client.

The IBM Spectrum Protect for Space Management client manages space on the following file systems:

- General Parallel File System (GPFS) on AIX® cluster
- GPFS on Linux x86_64 cluster
- GPFS on Linux on z Systems cluster

Related concepts:

Chapter 1, “Overview of the space management client,” on page 1

“Assignment of management classes to files” on page 24

“Registering the space management client with the IBM Spectrum Protect server” on page 17

General installation prerequisites and considerations

Before installing the space management client on any system, review the general requirements and considerations.

Installation prerequisites

You must install the IBM Spectrum Protect backup-archive client and you must have authority to install the space management client.

- You must have root user authority to install, set up, and use the space management client on your workstation.
- You must install and set up the IBM Spectrum Protect backup-archive client before you install the space management client. Both the space management client and IBM Spectrum Protect backup-archive clients share common code, the same options files, communication protocols, node registration, and storage.

Pre-existing files in the installation directory can be deleted

The `/usr/tivoli/tsm/client/hsm/bin` directory is one of the base directories into which the space management client product is installed. Any files that you place in this directory might be deleted during installation. Do not place the following files into this directory:

- `dsm.opt` files
- `dsm.sys` files
- Include-exclude files
- User-created files

Installation steps overview

The following steps are an overview of the installation procedure.

1. Before installing anything, read all of these general requirements. Read your system installation overview, specific environment requirements, and each installation step to ensure that you are ready to install the space management client.
2. Follow your system-specific installation procedure.
3. When you complete the installation, register your workstation as a node with the IBM Spectrum Protect server.
4. Modify the `dsm.sys` and `dsm.opt` configuration files. Both files are in the following directory:
 - For AIX: `/usr/tivoli/tsm/client/ba/bin`
 - For Linux: `/opt/tivoli/tsm/client/ba/bin`

Reinstallation or upgrade overview

Before you reinstall the space management client, stop all activity and do not access any files on file systems to which you added space management. The installation process fails otherwise. Follow the upgrade procedure as described for your operating system. If the steps tell you to uninstall and reinstall the backup archive client and API, you must complete those steps. Uninstalling and reinstalling brings those packages to the same level as the space management client.

You must not upgrade with a command like `rpm -U`. You must not refresh with a command like `rpm -F`. You must uninstall the old version and then install the new version.

The space management client limitations and considerations

The space management client has limitations.

- Journal-based backup is not supported when the space management client for AIX GPFS systems is installed.
- The space management client for GPFS systems manages only file systems belonging to the local (home) GPFS cluster; it does not manage remotely mounted file systems.
- In a GPFS environment, a small file that is less than the size of a GPFS block size can become larger after an HSM migration. GPFS adds meta information to the file during the migration. Because another block on the file system is allocated for the meta information, the space allocated for the file increases. If a file system is filled to its maximum capacity with many small files, it can run out of space during the file migration.
- On a GPFS cluster with Linux and AIX nodes, install the space management client on Linux nodes only or on AIX nodes only. Do not install the space management client on Linux nodes and AIX nodes that are on the same GPFS cluster.
- On a GPFS cluster with Windows systems and UNIX systems nodes, you must run backup-archive client processes on the same node that runs HSM processes.
- The space management client is not compatible with DMAPI applications that use any of the following DMAPI events on the space managed file system:
 - `DM_EVENT_READ`
 - `DM_EVENT_WRITE`
 - `DM_EVENT_TRUNCATE`
 - `DM_EVENT_PREUNMOUNT`

- DM_EVENT_NOSPACE
- DM_EVENT_DESTROY

Any DMAPI application on the space managed file system must respond to the DM_EVENT_MOUNT event with DM_RESP_DONTCARE.

Related concepts:

Chapter 3, “Configuring the space management client,” on page 19

“Space requirement for HSM control files” on page 188

Related information:

-  Installing the backup-archive clients
-  Journal-based backup

Installation overview for the space management client for AIX GPFS systems

Before you install the space management client on AIX General Parallel File Systems (GPFS) systems, review both the general and the system-specific requirements. If you are installing the product for the first time, use the steps for an initial installation, otherwise use the steps for an upgrade.

There are several installation limitations for the space management client for AIX GPFS systems:

- On AIX 6.1 and 7.1, the space management client can be installed in the global partition and supports transparent recall for both global and local workstation partitions (WPARs). Using HSM commands from a local WPAR is not supported. You cannot install the space management client in a local WPAR.

When you install the space management client on GPFS file systems, the installation process does the following tasks:

- Stops any space management daemons that are running.
- Removes any statement from the /etc/inittab file that loads the **dsmwatchd** command at system startup.
- Removes any statement from the /var/mmfs/etc/gpfsready script file that loads the other space management daemons at GPFS system startup.
- Extracts the HSM modules.
- Adds a statement to the /etc/inittab file that loads the **dsmwatchd** daemon at system startup.
- Adds a statement to the /var/mmfs/etc/gpfsready script file that loads the other space management daemons at GPFS system startup.
- Starts the space management daemons.

Table 1 indicates the packages available on the installation media in the /usr/sys/inst.images directory:

Table 1. The space management client for AIX GPFS systems installation packages

Package	Installs	Into this directory
tivoli.tsm.client.ba64.gpfs	The backup-archive client for AIX GPFS	/usr/tivoli/tsm/client/ba/bin
tivoli.tsm.client.hsm.gpfs	The space management client for AIX GPFS	/usr/tivoli/tsm/client/hsm/bin
tivoli.tsm.client.api.64bit	The API for AIX	/usr/tivoli/tsm/client/api/bin

For an initial installation, follow these steps:

1. If you want the GPFS policy engine to control automatic migration, you can disable the **dsmonitord** and **dsmscoutd** automatic migration daemons. Disabling these daemons conserves system resources. To disable the automatic migration daemons, start this command in a shell:

```
export HSMINSTALLMODE=SCOUTFREE
```

For information about configuring GPFS integration with the space management client, see Technote 7018848.

2. Edit the `dsm.opt` and `dsm.sys` files that are installed with the backup-archive client to configure the space management client.
3. Install the space management client on each node. For AIX clients, see “Installing the space management client for AIX systems” on page 9. For Linux clients, see “Installing the space management client for Linux GPFS systems” on page 14.
4. Make sure that after installation, the **dsmscald** daemon is running on at least one node.
5. Enable the Data Management Application Programming Interface (DMAPI) for GPFS for all file systems to which you plan to add space management. Enable DMAPI only once for each file system.
 - a. Unmount all GPFS file systems on all nodes within the GPFS cluster to which you plan to add space management.
 - b. Activate DMAPI management for the GPFS file systems with the following command: `mmchfs device -z yes`.

For information about GPFS commands and GPFS requirements for the IBM Spectrum Protect for Space Management client, go to the General Parallel File Systems product information and see `mmbackup` command: IBM Spectrum Protect requirements.
 - c. Remount all GPFS file systems on all nodes within the GPFS cluster.

The HSM daemons detect the initial state of each node and assign all nodes an instance number in relation to the GPFS cluster definition.
6. On the HSM owner nodes, add space management to each GPFS file system with the **dsmmigfs** command.
7. Use the **dsmmigfs enablefailover** command to enable failover of space management on the owner and source cluster nodes that participate in the failover group.

Related tasks:

“Editing the options files by using the backup-archive client” on page 20

Related reference:

“Adding space management to workload partitions on AIX V6.1 and V7.1 operating systems” on page 39

“**dsmmigfs add** and **update**” on page 128

“**dsmmigfs enablefailover**, and **disablefailover**” on page 143

Related information:

 [Journal-based backup](#)

Environment requirements for the space management client for AIX GPFS systems

Hardware, software, and communication requirements for HSM on AIX systems are listed in tech notes.

Environment requirements for HSM on AIX systems are listed in Technote 1248419.

Installing the space management client for AIX systems

You can install the space management client for AIX systems from a local directory where you copy the client files.

Before you begin

You must be logged in as the root user to install the product.

This installation procedure can be used to install new distributions or updates from downloaded installation media. The downloaded files that you use to install the space management client might be compressed. Depending on the package file format, either copy or extract the files to disk and use these instructions to install the components.

You can download the appropriate package file from one of the following websites:

- Download the space management client package from Passport Advantage or Fix Central.
- For the latest information, updates, and maintenance fixes, go to the IBM Support Portal.

About this task

Follow these steps from the AIX command line. These steps install the space management client for AIX General Parallel File System (GPFS) systems.

Procedure

1. Optional: Specify **HSMINSTALLMODE=SCOUTFREE**.

If you want the GPFS policy engine to control automatic migration, you can disable the **dsmonitord** and **dsmscoutd** automatic migration daemons. Disabling these daemons conserves system resources. To disable the automatic migration daemons, start this command in a shell:

```
export HSMINSTALLMODE=SCOUTFREE
```

For information about configuring GPFS integration with the space management client, see Technote 7018848.

2. Issue the **smitty install** command. If you are installing from a hard disk directory, remove the **.toc** file from that directory before you issue the **smitty** command.
3. Select **Install and Update Software** and press Enter.
4. Select **Install and Update from ALL Available Software** and press Enter.
5. In the **INPUT device / directory for software** field, select the following option and press Enter.
 - **/usr/sys/inst.images**
6. In the **SOFTWARE to Install** field press F4.

7. Scroll through the file sets and press F7 to select the IBM Spectrum Protect file sets that you want to install and press Enter. The space management client installation requires the following file sets:
 - IBM Spectrum Protect application programming interface
 - IBM Spectrum Protect backup-archive client common files
 - IBM Spectrum Protect backup-archive client base files
8. Select the options that you want and press Enter to begin the installation.

What to do next

After installing the space management client, you can register your workstation with the IBM Spectrum Protect server.

Related concepts:

“Registering the space management client with the IBM Spectrum Protect server” on page 17

Upgrading the space management client for AIX GPFS systems

You must remove the space management client for GPFS before installing a new version. You must deactivate the space management client, disable failover, and remove the space management client from all nodes.

About this task

If you want to upgrade the space management client, you must uninstall and install the backup-archive client and the API to get the product to the same version and level. A mixed installation is not supported. Activity on the space-managed file systems should be avoided during software installation. Access to migrated files during the upgrade is not possible, because the system would be suspended while trying to read a file stub.

To upgrade the space management client for GPFS, complete the following steps:

Procedure

1. Ensure that all HSM nodes (owner and backup) are in a consistent state and that all space-managed file systems are mounted on all HSM nodes.
2. Globally deactivate HSM on every node by issuing the command: **dsmmigfs globaldeactivate**.
3. Disable failover on every node by issuing the command: **dsmmigfs disablefailover**.
4. To figure out which node owns which cluster file system, issue the **dsmmigfs q -d** command on one of the nodes.
5. Remove the space management client from all nodes. You must not upgrade with a command like **rpm -U**. You must uninstall the old version and then install the new version.
6. Install the space management client on all nodes.
7. Globally reactivate HSM on every node by issuing the command: **dsmmigfs globalreactivate**.
8. Enable failover on every node by issuing the command: **dsmmigfs enablefailover**.

9. Take over each file system on its owner node as known from step 4 on page 10. With distributed HSM, the HSM daemons also run on these nodes, which are designated for the failover feature.

Related concepts:

“Installation overview for the space management client for AIX GPFS systems” on page 7

Related reference:

“**dsmmigfs globaldeactivate** and **globalreactivate**” on page 138

“**dsmmigfs enablefailover**, and **disablefailover**” on page 143

“**dsmmigfs deactivate**, **reactivate**, and **remove**” on page 137

“**dsmmigfs query**” on page 139

Uninstalling the space management client for AIX GPFS systems

Complete this procedure to uninstall the space management client for AIX GPFS systems.

About this task

Before you uninstall the space management client, read the space management client for AIX GPFS systems installation overview (“Installation overview for the space management client for AIX GPFS systems” on page 7).

If you do not remove HSM support from all managed file systems, the data of migrated files is not accessible after you remove the space management client.

Follow these steps to uninstall the space management client for AIX systems:

Procedure

1. To remove HSM support from all managed file systems, enter the **dsmmigfs remove** command.
2. Issue the **smitty remove** command.
3. In the **SOFTWARE name** field, press F4 to list the IBM Spectrum Protect file sets that you want to uninstall.
4. Select the IBM Spectrum Protect file sets that you want to uninstall and press Enter.
5. In the **PREVIEW only?** field, select N0 and press Enter.

Related concepts:

“Installation overview for the space management client for AIX GPFS systems” on page 7

Related reference:

“**dsmmigfs deactivate**, **reactivate**, and **remove**” on page 137

Installation overview for the space management client for Linux GPFS systems

Before you install the space management client on Linux General Parallel File System (GPFS) systems, review both the general and the system-specific requirements. If you are installing the product for the first time, use the steps for an initial installation. Otherwise, use the steps for an upgrade.

Note:

- HSM cluster installations are certified on IBM Linux Cluster 1350. See the IBM Redbooks®: *Linux Clustering with CSM and GPFS*.
- Also, see the recommendations that are provided with the IBM GPFS for Linux systems.

When you install the space management client on Linux GPFS file systems, the installation process makes the following changes:

- Stops any space management daemons that are running.
- Removes any statement from the `/etc/inittab` file that starts the **dsmwatchd** daemon at system startup.
- Removes any statement from the `/var/mmfs/etc/gpfsready` script file that loads the other space management daemons at GPFS system startup.
- Extracts the HSM modules.
- Adds a statement to the `/etc/inittab` file that starts the **dsmwatchd** daemon at system startup.
- Adds a statement to the `/var/mmfs/etc/gpfsready` script file that starts the other space management daemons at GPFS system startup.
- Starts the space management daemons.

Table 2 lists the packages available on the installation media for Linux on x86_64 systems:

Table 2. Linux x86_64 GPFS available packages

Package	Installs	Into this directory
TIVsm-API64.x86_64.rpm	The API for Linux x86_64 (64-bit only)	/opt/tivoli/tsm/client/api/bin64
TIVsm-BA.x86_64.rpm	The IBM Spectrum Protect backup-archive client (command line), the administrative client (command line), and the web backup-archive client (64-bit only)	/opt/tivoli/tsm/client/ba/bin
TIVsm-HSM.x86_64.rpm	The space management client for Linux x86_64 (64-bit only)	/opt/tivoli/tsm/client/hsm/bin

Table 3 lists the packages available on the installation media for Linux on z Systems:

Table 3. Linux z Systems GPFS available packages

Package	Installs	Into this directory
TIVsm-API64.s390x.rpm	The API for Linux on z Systems (64-bit only)	/opt/tivoli/tsm/client/api/bin64

Table 3. Linux z Systems GPFS available packages (continued)

Package	Installs	Into this directory
TIVsm-BA.s390x.rpm	The IBM Spectrum Protect backup-archive client (command line), the administrative client (command line), and the web backup-archive client (64-bit only)	/opt/tivoli/tsm/client/ba/bin
TIVsm-HSM.s390x.rpm	The space management client for Linux on z Systems (64-bit only)	/opt/tivoli/tsm/client/hsm/bin

For an initial installation, complete the following steps:

1. If you want the GPFS policy engine to control automatic migration, you can disable the **dsmonitord** and **dsmscoutd** automatic migration daemons. Disabling these daemons conserves system resources. To disable the automatic migration daemons, start this command in a shell:


```
export HSMINSTALLMODE=SCOUTFREE
```

 For information about configuring GPFS integration with the space management client, see Technote 7018848.
2. Edit the `dsm.opt` and `dsm.sys` files that are installed with the backup-archive client to configure the space management client.
3. Install the space management client on each node. For AIX clients, see “Installing the space management client for AIX systems” on page 9. For Linux clients, see “Installing the space management client for Linux GPFS systems” on page 14.
4. Make sure that after installation, the **dsmscald** daemon is running on at least one node.
5. Enable the Data Management Application Programming Interface (DMAPI) for GPFS for all file systems to which you plan to add space management. Enable DMAPI only once for each file system.
 - a. Unmount all GPFS file systems on all nodes within the GPFS cluster to which you plan to add space management.
 - b. Activate DMAPI management for the GPFS file systems with the following command: `mmchfs device -z yes`.

For information about GPFS commands and GPFS requirements for the IBM Spectrum Protect for Space Management client, go to the General Parallel File Systems product information and see `mmbackup` command: IBM Spectrum Protect requirements.
 - c. Remount all GPFS file systems on all nodes within the GPFS cluster.

The HSM daemons detect the initial state of each node and assign all nodes an instance number in relation to the GPFS cluster definition.
6. On the HSM owner nodes, add space management to each GPFS file system with the **dsmmigfs** command.
7. Use the **dsmmigfs enablefailover** command to enable failover of space management on the owner and source cluster nodes that participate in the failover group.

Related tasks:

“Editing the options files by using the backup-archive client” on page 20

Related reference:

“**dsmmigfs add** and **update**” on page 128

“**dsmmigfs enablefailover**, and **disablefailover**” on page 143

Environment requirements for the space management client for Linux GPFS systems

Hardware, software, and communication requirements for HSM on Linux systems are listed in tech notes.

Software requirements

Environment requirements for HSM on Linux on x86_64 systems are listed in Technote 1248771 .

Environment requirements for HSM on Linux on z Systems are listed in Technote 1966164.

Installing the space management client for Linux GPFS systems

You can install the space management client on Linux General Parallel File System (GPFS) systems from the product installation media.

Before you begin

You must be logged in as the root user to install the product.

This installation procedure can be used to install new distributions or updates from downloaded installation media. The downloaded files that you use to install the space management client might be compressed. Depending on the package file format, either copy or extract the files to disk and use these instructions to install the components.

You can download the appropriate package file from one of the following websites:

- Download the space management client package from Passport Advantage or Fix Central.
- For the latest information, updates, and maintenance fixes, go to the IBM Support Portal.

Procedure

1. Change directory to the installation packages directory.
2. Optional: Specify **HSMINSTALLMODE=SCOUTFREE**.

If you want the GPFS policy engine to control automatic migration, you can disable the **dsmonitord** and **dsmscoutd** automatic migration daemons. Disabling these daemons conserves system resources. To disable the automatic migration daemons, start this command in a shell:

```
export HSMINSTALLMODE=SCOUTFREE
```

For information about configuring GPFS integration with the space management client, see Technote 7018848.

3. Install the IBM Spectrum Protect clients in the order that is presented in Table 4 on page 15. During installation, these packages are installed in unique directories. Press **y** to each question when prompted.

Table 4. Package names and descriptions

To install	Issue this command
API	Install the API files in the /opt/tivoli/tsm/client/api/bin64 directory: <ul style="list-style-type: none"> On x86_64 systems: rpm -i TIVsm-API64.x86_64.rpm On z Systems: rpm -i TIVsm-API64.s390x.rpm
Backup-archive client (CLI and GUI) web client Administrative client	Install the backup-archive client (command line, GUI, and web client) in the /opt/tivoli/tsm/client/ba/bin directory: <ul style="list-style-type: none"> On x86_64 systems: rpm -i TIVsm-BA.x86_64.rpm On z Systems: rpm -i TIVsm-BA.s390x.rpm To circumvent the dependence check, you can use the --nodeps option, but then you must check the dependencies manually.
The space management client	Install the space management client commands and the space management agent in /opt/tivoli/tsm/client/hsm/bin directory: <ul style="list-style-type: none"> On x86_64 systems: rpm -i TIVsm-HSM.x86_64.rpm On z Systems: rpm -i TIVsm-HSM.s390x.rpm

Related reference:

“`hsmdisableautomigdaemons`” on page 94

Upgrading the space management client for Linux GPFS systems

You must remove the space management client for GPFS before installing a new version. You must deactivate the space management client, disable failover, and remove the space management client from all nodes.

About this task

If you want to upgrade the space management client, you must uninstall and install the backup-archive client and the API to get the product to the same version and level. A mixed installation is not supported. Activity on the space-managed file systems should be avoided during software installation. Access to migrated files during the upgrade is not possible, because the system would be suspended while trying to read a file stub.

To upgrade the space management client for GPFS, complete the following steps:

Procedure

1. Ensure that all HSM nodes (owner and backup) are in a consistent state and that all space-managed file systems are mounted on all HSM nodes.

2. Globally deactivate HSM on every node by issuing the command: **dsmmigfs globaldeactivate**.
3. Disable failover on every node by issuing the command: **dsmmigfs disablefailover**.
4. To figure out which node owns which cluster file system, issue the **dsmmigfs q -d** command on one of the nodes.
5. Remove the space management client from all nodes. You must not upgrade with a command like **rpm -U**. You must uninstall the old version and then install the new version.
6. Install the space management client on all nodes.
7. Globally reactivate HSM on every node by issuing the command: **dsmmigfs globalreactivate**.
8. Enable failover on every node by issuing the command: **dsmmigfs enablefailover**.
9. Take over each file system on its owner node as known from step 4. With distributed HSM, the HSM daemons also run on these nodes, which are designated for the failover feature.

Related concepts:

“Installation overview for the space management client for Linux GPFS systems” on page 12

Related reference:

“**dsmmigfs globaldeactivate** and **globalreactivate**” on page 138

“**dsmmigfs enablefailover**, and **disablefailover**” on page 143

“**dsmmigfs query**” on page 139

“**dsmmigfs deactivate**, **reactivate**, and **remove**” on page 137

Uninstalling the space management client for Linux GPFS systems

You must follow a specific procedure to uninstall the space management client for Linux GPFS systems.

About this task

To uninstall the space management client for Linux GPFS systems, complete the following steps.

Procedure

1. Remove HSM support from all managed file systems, by issuing this command: **dsmmigfs remove**. If you do not remove HSM support from all managed file systems, the data of migrated files will not be accessible after you remove the HSM package.
2. Issue the command, **rpm -e TIVsm-HSM**

What to do next

After uninstalling, you can install the latest version.

Related concepts:

“Installation overview for the space management client for Linux GPFS systems” on page 12

Related reference:

Registering the space management client with the IBM Spectrum Protect server

After installing the space management client, your system must be registered as a client node with the IBM Spectrum Protect server before you can request services from that server. Your IBM Spectrum Protect server administrator has set the registration either to closed or to open.

Registering your node by using open registration

In open registration, the root user registers your workstation as a client node with the IBM Spectrum Protect server. Your workstation must be registered before anyone can use an IBM Spectrum Protect client on that node.

About this task

Open registration has the following features:

- Your client node is assigned to a policy domain named STANDARD.
- The root user can set the appropriate value for the **compression** option in your `dsm.sys` file.
- The root user can delete archived copies, but not backup versions, of files from IBM Spectrum Protect storage. Users can delete archived files that they own.

Note: Your IBM Spectrum Protect administrator can change these defaults at any time.

Follow these steps to register your workstation with the IBM Spectrum Protect server:

Procedure

1. Start a session with the command-line interface by issuing the **dsmc** command.
2. When you are prompted for information to register your workstation with a server that is identified in your `dsm.sys` file, supply the following information:
 - If a password is required, supply a password.
 - Supply contact information, such as your name, user ID, and telephone number
3. To register your workstation with more servers, issue the **dsmc** command with the **servername** option for each server. For example, where `dsmserve` is the name of a server that you identified in your `dsm.sys` file, issue the command:

```
dsmc -servername=dsmserve
```
4. Enter information at the prompts to register your workstation with the server that you specify.

Closed registration

With closed registration, you must provide your IBM Spectrum Protect administrator with information to register your workstation as a client node with the server.

If your enterprise uses closed registration, your IBM Spectrum Protect administrator registers your node. Your IBM Spectrum Protect administrator requires the following information:

1. Your node name: the value that the **hostname** command returns or the node name that you specified by using the **nodename** option
2. The initial password that you want to use, if a password is required
3. Contact information, such as your name, user ID, and telephone number

Your administrator defines the following policies for your node:

- The policy domain to which your client node belongs

Note: A policy domain contains policy sets and management classes that control how IBM Spectrum Protect manages the files you back up, archive, or migrate.

- Whether you can compress files before you send them to the server
- Whether you can delete backup and archive data from IBM Spectrum Protect storage

Password management

The backup-archive client and the space management client use the same password when both clients contact the same IBM Spectrum Protect server. One password is required for each IBM Spectrum Protect server that your client node contacts for services.

If you specify a migration server by using the **migrateserver** option in your `dsm.sys` file, the password that you set applies to the migration server.

If you specify a default server and you do not specify a migration server, the password that you set applies to the default server.

If you do not specify either a migration server or a default server, the password that you set applies to the named server in the first stanza of your `dsm.sys` file.

Note: Your IBM Spectrum Protect password can be changed only from the command line. To change that password from the command line, enter the command (where `oldpw` is the old password and `newpw` is your new password):

```
dsmsetpw oldpw newpw
```

You can set the **passwordaccess** option to generate in your `dsm.sys` file. IBM Spectrum Protect encrypts and stores your password locally and automatically generates a new password for your client node each time it expires. You are not prompted for the password. To receive a note with the password when a new password is generated, set the `mailprog` option in your `dsm.sys` file.

Related reference:

“**dsmsetpw**” on page 165

Chapter 3. Configuring the space management client

Before you configure space management for a file system, you must configure the space management client itself. You also must configure some options that apply to all space-managed file systems.

The space management client is configured by setting options in the `dsm.sys` and `dsm.opt` files. The options generally affect all space-managed file systems, and determine the following things:

- The IBM Spectrum Protect servers to which your files migrate or premigrate
- The include-exclude file that applies for each IBM Spectrum Protect server.
- How often space usage is checked on your file systems
- How often your file systems are automatically reconciled
- How often candidates are searched for automatic migrations
- How many automatic migration processes for each file system can migrate files in parallel
- How many days to keep an obsolete file copy in IBM Spectrum Protect storage before the copy expires.
- Whether to use two-way orphan check reconciliation (for GPFS only)

During installation of the IBM Spectrum Protect for AIX and Linux backup-archive clients, sample options files (`dsm.sys.smp` and `dsm.opt.smp`) are placed in the following directories.

- On AIX: `/usr/tivoli/tsm/client/ba/bin64`
- On Linux: `/opt/tivoli/tsm/client/ba/bin`

If you install both the backup-archive client and the space management client at the same time, copy and rename the sample options files. Modify them for both clients. If you previously installed the backup-archive client and you set up your options files, modify them for the HSM client.

After you select options for the space management client, you must restart all space management daemons to activate the changes. All space-managed file systems must be mounted either automatically or manually each time you restart your system to start space management.

The space management client shares the following common files and code with the AIX and Linux backup-archive clients:

- Communication protocols
- The `dsm.opt` options file
- The `dsm.sys` options file
- The include-exclude file
- Node registration
- Server file spaces

On GPFS file systems, you can use the GPFS policy engine to monitor space thresholds and search for migration candidates.

For information about configuring GPFS integration with the space management client, see Technote 7018848.

Related concepts:

Chapter 11, “Options files reference,” on page 89

Related information:

 Backup-archive client processing options

Displaying space management client options

To display space management client options, use the **dsmmigquery** command with the **options** parameter.

Related reference:

“**dsmmigquery**” on page 147

Editing the options files by using the backup-archive client

The space management client shares the option files, `dsm.opt` and `dsm.sys`, with the IBM Spectrum Protect for AIX and Linux backup-archive clients. To edit the options files, you can use the Preferences Editor window of the backup-archive client GUI. You can also edit the files in a text editor.

About this task

Note: You can use the backup-archive GUI online help to get more detailed information by clicking the **Help** button on the Preferences Editor window.

Follow these steps to edit the options files from the backup-archive GUI:

Procedure

1. Start the backup-archive client GUI by issuing the command **dsmj**.
2. Open the Preferences Editor window by choosing **Edit > Preferences** on the menu.
3. Select the tab for the options that you want to edit and make any needed changes.

Editing the options file `dsm.sys`

The options that you define in the `dsm.sys` file for the space management client affect automatic migration, reconciliation, and recall.

About this task

You must have root user authority to set the options in your `dsm.sys` file.

In the `dsm.sys` file, group your options into stanzas for each server that your client node contacts for backup, archive, and space management services.

Options are processed following this order:

1. Options that are defined on the server with server-enforced client options (the client cannot override the value)
2. Options that are entered locally on the command line
3. Options that are defined on the server for a schedule
4. Options that you enter locally in your options file

5. Options that are received from the server with client options that the server does not enforce (the client can override the value)
6. Default option values

For options that you can set in the `dsm.sys` file, see Chapter 11, “Options files reference,” on page 89

Note: You can specify both a default server and a migration server in the `dsm.sys` file. Use the **defaultserver** option and the **migrateserver** option. If you do not specify migration server and default server, the server that you specify in the first stanza of the `dsm.sys` file becomes the default server. If you specify a migration server with the **migrateserver** option, it overrides the server that you specify with the **defaultserver** option.

Procedure

1. Obtain the server information from your IBM Spectrum Protect administrator.
2. Edit your `dsm.sys` file to include the server to which you want to connect for space management services.
3. Assign a name to the server that you want to contact for space management services. For each **servername** entry, include a **commmethod** entry to specify the communication method to use for client and server communications.
4. Issue a value for each option and remove the leading asterisk (*). You can specify options for more than one server.

Example

This code is an example of a `dsm.sys` server stanza.

```

DEFAULTServer          server1
MIGRATEServer         server2
CHECKThresholds       2
CANDIDATESInterval    12
MAXCANDprocs          5
RECOncileinterval    1
MAXRECOncileproc      5
MAXThresholdproc      5
MINMIGFILESIZE        8192
MIGFILEEXPIration    10
MINRECALLdaemons     5
MAXRECALLdaemons     15
CHECKFororphans       no
MAXMIGRators          1
KERNelmessages        no
OVERLAPRECALL         no
Servername      server1
  COMMmethod          TCPip
  TCPPort              1500
  TCPServeraddress    server3.almaden.ibm.com
  Passwordaccess      generate
  Mailprog             /usr/bin/xsend root
  Groups              system tsm
  Users               steiner chron wang nguyen
  Inclxcl              /adm/tsm/backup.excl
  ERRORProg           /bin/cat
Servername      server2
  COMMmethod          SNA1u6.2
  PARTner1uname       raptor
  TPname              appcdel
  CPICM0dename        appc
  Passwordaccess      generate
  Mailprog             /usr/bin/xsend root

```

Groups	system tsm
Users	sullivan tang stewart
Incl excl	/adm/tsm/migrate.excl
ERRORProg	/bin/cat

Editing the options file dsm.opt

Most of the options in the dsm.opt file affect the backup-archive client. However, some options affect the space management client.

About this task

Table 5 provides a brief description of the space management options that you can set in your dsm.opt file.

Note: The IBM Spectrum Protect **nfstimeout** option is ignored for the space management client. The space management client operates only on local file systems.

Procedure

1. Group the options into stanzas for each server that your client node contacts for backup, archive, and space management services.
2. Enter a value for each option and remove the leading asterisk (*).

Table 5. space management client options in the dsm.opt file

Option	Default	Description
defaultserver	The server that you identify in the first stanza of the dsm.sys file.	Specifies the name of the default IBM Spectrum Protect server to contact for space management services when you do not specify a server name on the migrateserver option. The value of defaultserver in the dsm.sys file overrides defaultserver in the dsm.opt file. See “ defaultserver ” on page 92.
hsmgroupedmigrate	NO	When this option is set to YES, the HSM transaction grouping is activated. Multiple files are migrated in each transaction with the IBM Spectrum Protect server. The files are migrated when either the transaction byte limit or the transaction group limit is reached. The transaction group limit is specified with the txngroupmax option, which is set on the IBM Spectrum Protect server. See “ hsmgroupedmigrate ” on page 97.
migrateserver	The value of defaultserver option.	Specifies the IBM Spectrum Protect server to which you want to migrate files from your client node. Specify only 1 server for each client node. The value of migrateserver in the dsm.sys file overrides migrateserver in the dsm.opt file. See “ migrateserver ” on page 106.
restoremigstate	YES	Restores a file to stubbed (migrated) state. The file must have been backed up after migration or premigration. IBM Spectrum Protect records the migration state of files during backup, so only files that were migrated or premigrated before backup can be restored to stubbed (migrated) state. See “ restoremigstate ” on page 109.

3. Set the **DSM_CONFIG** environment variable to point to your dsm.opt file.

Related reference:

“Environment variables” on page 28

Configuring the space management client to connect to a secondary IBM Spectrum Protect server

If the primary IBM Spectrum Protect server for the space management client is unavailable, you can manually configure the space management client to connect to a secondary server. You can recall files from the secondary IBM Spectrum Protect server but cannot migrate files to the secondary server.

Before you begin

The primary IBM Spectrum Protect server for the space management client must be one that replicates client node data.

About this task

The IBM Spectrum Protect server that the client connects to during normal production processes is called the *primary server*. When the primary server is set up for node replication, the data for client nodes can be replicated to the *secondary server*.

The backup-archive client can automatically fail over to the secondary server when it is configured for failover.

The space management client, however, does not automatically fail over to the secondary server. You must manually edit the `dsm.sys` file to connect to the secondary server. Any secondary server information in the **replservername** stanza and **myreplicationserver** option is ignored by the space management client.

Procedure

To configure the space management client to connect to the secondary server, complete the following steps:

1. Edit the `dsm.sys` file. Add a **servername** stanza that contains connection information for the secondary server. The following stanza is an example of a secondary server stanza:

```
Servername      lifeboat_server
  COMMmethod    TCPip
  TCPPort       1500
  TCPServeraddress  server4.almaden.ibm.com
  Passwordaccess  generate
  Mailprog      /usr/bin/xsend root
  Groups        system tsm
  Users         steiner chron wang nguyen
  Inclxcl       /adm/tsm/backup.excl
  ERRORProg     /bin/cat
```

2. Update the **defaultserver** or **migrateserver** options to point to the secondary server stanza.

```
DEFAULTServer      lifeboat_server
MIGRATEServer      lifeboat_server
```

3. For non-root users, edit the `dsm.opt` file and update for value in the **defaultserver** option to point to the secondary server stanza that you added in the `dsm.sys` file.

```
DEFAULTServer      lifeboat_server
```

4. Restart the space management client, including all space management daemons.

Results

Functions that require write access to the secondary IBM Spectrum Protect server are not available for a space managed file system. You cannot back up, archive, or migrate files to the secondary server from a space managed file system. You can restore, retrieve, and recall from the secondary server.

With the backup-archive client, you cannot restore migrated files as stub files. Migrated files can only be restored as resident files.

What to do next

To switch back to the primary IBM Spectrum Protect server, repeat the steps, but change the value of the **servername** option from the target server to the source server.

Related tasks:

“Stopping the space management daemons” on page 85

Related information:

 Automated client failover configuration

Optional setup features

You can specify which files are eligible for migration. You can use clustered file systems and LAN-free data transfer and set logging options.

Assignment of management classes to files

A management class indicates a storage policy for a file. The storage policy determines whether a file can be migrated and how the migrated copy is stored.

The IBM Spectrum Protect administrator defines management classes that contain specific requirements or policies for migrating files to storage. You assign these management classes to files on your local file systems. The management class that you assign to a file determines file eligibility for migration. Use the default management class for some or all of your files. Assign different management classes to specific files or groups of files with one or more include statements in your include-exclude options file.

A management class can contain a backup copy group and an archive copy group. Copy groups contain attributes that control the generation, destination, and expiration of backup versions of files and archived copies of files. Table 6 on page 25 lists the space management attributes and their defaults that might be included in a management class.

Table 6. The space management client attributes in a management class

Attribute	Default	Description
spacemgtechnique	None	Specifies that a file is eligible for automatic and selective migration, selective migration only, or neither. This attribute has the following values: Auto The file is eligible for both automatic and selective migration. Selective The file is eligible for selective migration only. None The file is not eligible for migration. Note: If you use the default management class named Standard that is shipped with IBM Spectrum Protect, and your administrator did not change the default setting for the spacemgtechnique attribute, files are not migrated from your workstation.
automignonuse	0	Specifies the number of days (0 through 9999) that must elapse since you last accessed the file before it is eligible for automatic migration.
migrequiresbkup	YES	Determines whether a current backup version of the file must exist on your migration server before the file is eligible for automatic or selective migration. The following values are valid: YES A current backup version must exist on the IBM Spectrum Protect server to which the file is migrated. NO A current backup version is not required. Note: If you set this attribute to YES, IBM Spectrum Protect checks for a current backup version of the file on your migration server only. If a current backup version does not exist on your migration server, the file is not migrated.
migdestination	spacemgpool	Specifies the name of the storage pool in which IBM Spectrum Protect stores the file when it migrates.

Table 7 identifies appropriate management classes for some tasks.

Table 7. Assigning management classes to your files

Task	Assignment
You want to migrate and back up a file.	Assign a management class to a file with space management attributes and backup-archive copy groups that you want to use for that file. Assign only one management class to a specific file.
Your client node communicates with the same server for both space management and backup-archive services.	Assign a management class that contains space management attributes and backup-archive copy groups that you assigned to the backup-archive client to use for that file.
You migrate files to one server and back up and archive files to one or more different servers:	<ul style="list-style-type: none"> Specify a different include-exclude options file for each server. Assign only one management class to a file in a specific include-exclude options file. Assign different management classes to files in different include-exclude options files. <p>For example, if you back up files in /home/holland to Server1, the include-exclude options file that you use for Server1 might specify a management class named <i>mgmt1a</i> for a file named /home/holland/testfile. This management class must contain an appropriate backup copy group for the file.</p> <p>If you migrate files in the /home file system to Server2, the include-exclude options file that you use for that server might specify a management class named <i>mgmt2b</i> for the same file. That management class must contain appropriate space management attributes for the file.</p>

Related information:

 [Storage management policies](#)

Displaying management class information

To display information about management classes, use the **dsmmigquery** command with the **-mgmtclass** parameter.

Related reference:

[“dsmmigquery” on page 147](#)

Related information:

 [Storage management policies](#)

Include-exclude file options

You can use an include-exclude options file to exclude or include specific files from space management and to assign specific management classes to these files. For example, you might want to keep certain files, on your local file system always, and thus need to exclude them from migration. Or, you might also want to include certain files for backup or migration.

Note:

- If you do not create an include-exclude options file, all files are considered for backup services and the default management class is used.
- On GPFS file systems, do not use the **incl excl** option in a failover environment. Unlike `dsm.opt` and `dsm.sys` files, include-exclude options files are not shared between different nodes of a failover group. Add the include-exclude list directly to the `dsm.sys` file or verify that the include-exclude list files match on all nodes that participate in the local failover group or node set.

For all other file systems, use the **incl excl** option in your `dsm.sys` file to specify the name of your include-exclude options file. You can create an include-exclude options file for each IBM Spectrum Protect server that your client node contacts for services. For example, if your `dsm.sys` file contains two stanzas with options for two IBM Spectrum Protect servers, you can include an **incl excl** option in each stanza. Each **incl excl** option can point to a different include-exclude options file. The files that you create must reside in a directory to which all users on your workstation have read access.

Follow these rules when you include or exclude files from either backup or migration:

- Use include statements to include files for backup or migration.
- Use exclude statements to exclude files or directories from backup or migration.
- End the specification for an include or exclude statement with a file name. You can use a wildcard in place of a specific file name.

The following options include or exclude objects from space management.

Table 8. Include and Exclude Statements

Option	Description
excl	Excludes a file or a group of files from space management.
<code>excl</code> .backup	Excludes a file from backup.

Table 8. Include and Exclude Statements (continued)

Option	Description
<code>exclude.file</code>	Excludes a file or a group of files from space management.
<code>exclude.file.spacemgmt</code>	Excludes a file from HSM services only. Use this option when you have both the backup-archive client and the HSM client installed.
<code>exclude.spacemgmt</code>	Excludes files and directories from HSM services only. Use this option when you have both the backup-archive client and the space management client installed.
include	Includes files for backup and HSM services or LAN-free data transfer.
<code>include.file</code>	Includes a file for backup.

The following example displays a sample include-exclude options file:

```
exclude /.../core
include /home/.../* personal_files
include /home/davehil/dsnew/.../*
include /home/davehil/driver5/.../* source_code
exclude.spacemgmt /home/jones/proj1/status/.../*
exclude /home/root/cron.log
```

In the sample include-exclude options file, *personal_files* and *source_code* identify management classes that are assigned to specific files. If you do not assign a management class to your files, the default management class is used.

Include-exclude options are processed from the bottom up. For example, when you build a migration candidates list for a file system, each file is tested against the options in the include-exclude options file beginning with the last option that you specified. If a match is found, it does not test the file against any additional options. It either excludes or includes the file as the option specifies. If a match is not found, the file is implicitly included for space management and backup services.

If you exclude a file from space management after it migrates to storage, it remains migrated until it is automatically or selectively recalled. After it is recalled, it is no longer eligible for migration. To prevent a file from migrating to storage from a local file system, assign the file to a management class without automatic or selective migration.

Related information:

-  Using IBM Spectrum Protect include and exclude options with the IBM Spectrum Scale mmbackup command, technote 1699569
-  Configuring IBM Spectrum Scale Active File Management

Creating an include-exclude list

The space management client shares the include-exclude list with the IBM Spectrum Protect for AIX and Linux backup-archive clients. You can use the backup-archive GUI or the command line to create and edit your include-exclude list.

About this task

Note:

- Before you use either method to create the include-exclude list, determine your include and exclude requirements.
- You can use the backup-archive client's online help to get more detailed information by clicking the **Help** on the Preferences Editor window.

To create or edit the include-exclude list by using the client backup-archive GUI, complete the following steps.

Procedure

1. Start the client backup-archive GUI by issue the command **dsmj**.
2. Start the preferences editor by choosing **Edit > Preferences** from the menu.
3. Select the **Include-Exclude** tab at the left side of the Preferences Editor window.
4. Click **Add** to add a new include-exclude statement. You can select a statement that you want to update, remove, or move up or down the statements list.

What to do next

To create an include-exclude list from the command line, complete the following steps.

1. Create an empty file in any directory to which all users on your workstation have read access.
2. Enter your include and exclude statements.
3. In the `dsm.sys` file, use the **incl excl** option to identify the file you created.

Environment variables

Environment variables determine where IBM Spectrum Protect clients create program files.

The following environment variables set the location of space management client program files. The environment variables cannot be set to the root directory.

Table 9. Environment Variables

Variable	Description
DSM_DIR	Points to the resource files, the <code>dsm.sys</code> file, and the executable file, <code>dsmt.ca</code> . If you do not set the DSM_DIR variable, the space management client searches for the executable files in the installation directory.
DSM_CONFIG	Points to your <code>dsm.opt</code> options file. <ul style="list-style-type: none">• If you do not set the DSM_CONFIG variable, the space management client searches for the options file in the directory to which the DSM_DIR variable points.• If you do not set the DSM_DIR variable, the space management client searches for the options file in the installation directory.

Table 9. Environment Variables (continued)

Variable	Description
DSM_LOG	<p>Points to the directory of the dsmerror.log log file. The error log file contains information about any errors that occur during processing.</p> <ul style="list-style-type: none">• If you define the DSM_DIR variable but you do not define the DSM_LOG variable, messages are written to the dsmerror.log log file in the directory that you specified in the DSM_DIR variable.• If you do not define the DSM_LOG variable or the DSM_DIR variable, error messages are written to the dsmerror.log log file in the current directory. You receive a warning message if IBM Spectrum Protect cannot write messages to the log file. Processing continues.

Bourne and Korn shell variables

To set up your Bourne or Korn shell, add environment variables to the .profile file in your \$HOME directory.

The following example shows HSM variable definitions:

```
DSM_DIR=/home/hsmuser  
DSM_CONFIG=/home/hsmuser/dsm.opt  
DSM_LOG=/home/hsmuser  
export DSM_DIR DSM_CONFIG DSM_LOG
```

- The DSM_DIR = /home/hsmuser entry identifies the path for the executable files, the resource files, and the client system options file.
- The DSM_CONFIG=/home/hsmuser/dsm.opt entry sets the path and file name for the client user options file.
- The DSM_LOG=/home/hsmuser entry sets the location of the dsmerror.log file.

C shell variables

To set up your C shell, add the **DSM_CONFIG** and **DSM_LOG** variables to the .cshrc file in your \$HOME directory.

The following example shows HSM variable definitions:

```
setenv DSM_CONFIG /home/hsmuser/dsm.opt  
setenv DSM_LOG /home/hsmuser
```

The /home/hsmuser/dsm.opt path identifies the path and file name for your client system options file. The /home/hsmuser path sets the location of the dsmerror.log file.

LAN-free data transfer for HSM

The space management client supports LAN-free data transfer, which shifts the movement of client data from the communications network to a storage area network (SAN). Shifting the client data movement from the communications network to a SAN decreases the load on the IBM Spectrum Protect server.

The SAN provides a path that allows migration and recall of data to and from a SAN-attached storage device. Client data moves over the SAN to the storage device via the IBM Spectrum Protect Storage Agent. The IBM Spectrum Protect Storage Agent must be installed on the same system as the client.

LAN-free prerequisites

Before you establish LAN-free support in your HSM environment, you must check the prerequisites.

To enable LAN-free support, the following prerequisites are required.

- The IBM Spectrum Protect Version 5.5 or later client and server is required.
- You must install and configure the IBM Spectrum Protect for SAN storage agent on the client workstation.

Related information:

 [Storage agent overview](#)

LAN-free options

After you install and configure the IBM Spectrum Protect Managed System for SAN feature on the client workstation, use options to enable LAN-Free data transfer.

Table 10 lists options that you can use to set up LAN-free data transfer.

Table 10. LAN-free data transfer options

Option	Description
<code>enablelanfree</code>	Specifies whether to enable an available LAN-free path to SAN-attached storage device. Set this option to YES.
<code>lanfreecommethod</code>	Specifies a communication protocol between the client and the Storage Agent. Set this option to TCPip.
<code>lanfreetcpport</code>	Specifies the TCP/IP port number where the Storage Agent is listening. Set this option to 1530.

You can use the **include** and **exclude** options to control LAN-free data transfer.

To include files:

Assuming that /hsm1 is a space-managed file system, include the files in the /hsm1/clientdata/lanfree directory for LAN-free data transfer and assign these files to a LAN-free enabled management class with the following command: `include /hsm1/clientdata/lanfree lanfreemgmtclass`

To exclude files:

Assuming that /hsm1 is a space-managed file system and the default management class is LAN-free, exclude the files in the /hsm1/clientdata directory from LAN-free data transfer with the following command:

```
exclude /hsm1/clientdata/*
```

Related information:

 [Backup-archive client processing options](#)

Configuring HSM on GPFS clusters

You can configure a space-managed file system to become a part of a resource group in a General Parallel File System (GPFS) cluster. If there is a system failure, you can access your data from another system.

About this task

The space management client on GPFS file systems also supports cascading and rotating takeover relationships in the same manner as the backup-archive client. The `dsm.opt` file and the `dsm.sys` file are in the standard `/opt/tivoli` directory for Linux GPFS, or the `/usr/tivoli` directory for AIX GPFS. You can also use the **DSM_CONFIG** environment variable that you set in the start script. However, the behavior must be identical on all nodes. Across different nodes, the server stanza must have the following characteristics:

- The server stanza must contain a node name, and the node name must be the same for the server stanza on each system.
- The node name can be a cluster name. The node name is not required to be the actual name of any node in the cluster.
- The server stanza must point to the same server on each system.

You can increase file transfer performance, migration, and recall capabilities to other GPFS nodes within a cluster environment. Use the **asnodename** option to share the GPFS file space on the server.

You can control distributed recall by setting the option **hsmdistributedrecall** to `NO` on all source nodes that you do not want to recall files.

To enable distributed migration and recall capabilities follow these steps:

Procedure

1. Install the space management client on all participating GPFS nodes in the cluster environment.
2. Register each space management client node with the IBM Spectrum Protect server. These nodes are used for managing HSM file systems and are used for distributed HSM. Nodes that are used only to access data of space-managed file systems do not require HSM. Nodes that are used only to access data of space-managed file systems are not required to be registered on the IBM Spectrum Protect server.
3. The IBM Spectrum Protect server administrator must use the **grant proxynode** server command. Use the command to grant proxy authority to the source nodes to access the node name where the GPFS file systems are stored.
4. Update the options file on each source node. Specify the **asnodename** option to access the common file space for the space-managed file systems on the IBM Spectrum Protect server.

Example

Configuration examples are available in the IBM Spectrum Protect Space Management-GPFS integration field guide.

For information about configuring GPFS integration with the space management client, see Technote 7018848.

Related concepts:

“Installation overview for the space management client for Linux GPFS systems” on page 12

“Installation overview for the space management client for AIX GPFS systems” on page 7

Related reference:

“[hsmdistributedrecall](#)” on page 95

Related information:

 [asnodedname](#) client option

 [GRANT PROXYNODE](#) administrative command

HSM limitations and requirements on GPFS file systems

Space management has the following limitations and requirements on General Parallel File System (GPFS):

- The management class information is for the default migration server only.
- The server options information is for the default migration server only.
- Every space management node must run the same HSM version.
- The backup-archive client cannot restore stub files to a GPFS file system that has more storage pools than the default system storage pool. Stub files are restored to their resident state. GPFS stores the GPFS pool ID in extended attributes. The backup-archive client cannot store these extended attributes independent from the file content.

HSM support for GPFS file systems is not integrated with the backup-archive client support. For example, the space management client refers to the *file_system/.SpaceMan/hsmfsconfig.xml* file to determine which IBM Spectrum Protect server to contact. The space management client might contact a different server for each file system. In contrast, the backup-archive client uses other methods to determine which server to use for a backup process.

HSM can be configured for GPFS file systems that use Active File Management (AFM). For information about running HSM on GPFS file systems that use AFM, see [Configuring IBM Spectrum Scale Active File Management](#).

Related information:

 [mmbackup](#) command: IBM Spectrum Protect requirements

 [Using IBM Spectrum Protect include and exclude options with the IBM Spectrum Scale mmbackup command, technote 1699569](#)

Logs for HSM activity and error messages

The space management client logs HSM activity and error messages. An HSM log contains information about file migration and recall, threshold migration, reconciliation, and starting and stopping the HSM daemon. An error log records error messages.

Logs for HSM activity

You can analyze an HSM log to determine the current state of the system. For instance, the HSM logs can indicate when a recall starts but does not finish within the last hour. The administrator can analyze a particular recall and react appropriately.

In addition, an administrator might analyze an HSM log to optimize HSM usage. For example, if the HSM log indicates that 1000 files are recalled at the same time, the administrator can analyze the log records. The administrator might suggest that the files can be first compressed into an archive file and then migrated.

To configure the HSM log file, set the following options in the `dsm.sys` file:

- **hsmlogeventflags**
- **hsmlogmax**
- **hsmlogname**
- **hsmlogretention**
- **hsmlogsampleinterval**

You can briefly change the HSM log file name if you specify the **logname** parameter for the following commands. The new log captures the results of only the operation that was started by the command.

- **dsmmigrate**
- **dsmrecall**
- **dsmmigundelete**
- **dsmmigfs**
- **dsmdf**

Daemons and commands that do not specify the **logname** parameter write log entries as follows:

- Log entries are written to the file specified with the **hsmlogname** option in the `dsm.sys` options file.
- If **hsmlogname** is not specified, log entries are written to the default log file.

Logs for error messages

To configure the error log file, set the following options in the `dsm.sys` file:

- **errorlogmax**
- **errorlogname**
- **errorlogretention**

You can briefly change the error log file name if you specify the **errorlogname** parameter on any space management client command. The new log captures the errors of only the operation that was started by the command.

Related reference:

- “**hsmlogeventflags**” on page 98
- “**hsmlogmax**” on page 99
- “**hsmlogname**” on page 99
- “**hsmlogretention**” on page 100
- “**hsmlogsampleinterval**” on page 101
- “**errorlogname**” on page 93
- “**dsmmigrate**” on page 149
- “**dsmrecall**” on page 157
- “**dsmmigundelete**” on page 153
- “**dsmdf**” on page 123
- “**dsmmigfs query**” on page 139

Related information:

 [errorlogmax client option](#)

 [errorlogretention client option](#)

Chapter 4. Adding and configuring space management for a file system

When you add space management to a file system, you indicate how and when files are migrated and recalled. You can also deactivate, reactivate, and remove space management from a file system.

Note:

- To add and configure space management for a file system, you must have root user authority.
- The space management client can work only on mounted file systems. Before you can migrate files to IBM Spectrum Protect server storage, first mount your file systems either automatically or manually and then add space management.
- During the mount process and while the space management client is adding space management to your file systems, do not attempt to access any files in your file systems. Do not perform any tasks against your file systems.

When you add space management to a file system, the space management client completes the following tasks:

- Creates a hidden directory for the file system named `.SpaceMan` that stores certain information objects that are required for space management.
- Creates the `hsmfsconfig.xml` file in the `.SpaceMan` directory of the file system. This file contains the space management settings that you selected.

Remember: The `.SpaceMan` directory is not processed by the IBM Spectrum Protect backup-archive client. Copy the `hsmfsconfig.xml` file to a directory that is included for automatic backups so that the IBM Spectrum Protect backup-archive client can process the file.

- Updates mount information for the native file system.
- Starts space management for the file system.

For GPFS file systems, the `dsmwatchd` daemon starts at system startup with an entry in the `etc/inittab` file or with the `initctl` service. Ensure that DMAPI is enabled on all GPFS file systems that the space management client manages. Issue the following command to query this information: `/usr/lpp/mmfs/bin/mm1sfs DevicePath -z`.

If Data Management Application Programming Interface (DMAPI) is disabled, enable it with following command: `/usr/lpp/mmfs/bin/mmchfs DevicePath -z yes`.

On GPFS file systems, you can change the value for the DMAPI enablement to YES only if the file system is unmounted on all nodes of the cluster. When DMAPI is enabled, the file system can be mounted only if a `dsmrecalld` daemon is set up on one of the cluster nodes within the GPFS cluster.

For information about configuring GPFS integration with the space management client, see Technote 7018848.

HSM-created stub files on your space-managed file systems are bound to the space-managed file system. You cannot do the following tasks:

- Move stub files as native stub files (without the migrated data) to other file systems, even if the other file system is space-managed.

- Use the IBM Spectrum Protect raw device backup utility or any other image-based backup utility to restore a space-managed file system to a system other than the original system.
- Use space-managed file systems within cluster replication tools.
- Use any other tools that transfer images between systems.

The following are more considerations:

- You cannot add space management to file systems such as the root file system and the temp file system.
- You can add space management to a nested file system.
- You can add space management to an exported file system.

Do not add space management to the /usr and /var file systems. All of those file systems contain files that your operating system uses regularly.

Related tasks:

“Adding space management to nested file systems”

“Adding space management to an exported file system” on page 37

Related reference:

Appendix A, “Control files in the .SpaceMan directory,” on page 187

Adding space management to file systems

To add space management to a file system , use the **dsmmigfs** command with the **-add** parameter.

Related reference:

“**dsmmigfs add** and **update**” on page 128

Adding space management to nested file systems

You can add space management to a nested file system.

About this task

A nested file system is a file system mount point that is contained within another file system. For example:

```
/test /test/migfs1
```

The /test file system is a parent file system and /test/migfs1 is a nested file system within /test. They are both mount points.

Procedure

1. Unmount the nested file systems.
2. Add space management to the parent file system.
3. Remount the nested file systems.
4. Optional: Add space management to the nested file systems. You can add space management to any, all, or none of the nested file systems.

Related tasks:

“Adding space management to file systems”

Adding space management to an exported file system

HSM supports only the NFS (Network File System) interface to export a file system.

About this task

To add space management to a file system that the NFS server exports, follow these steps:

Procedure

1. Instruct all NFS clients to unmount the exported file system.
2. To view which clients mounted the exported file system, issue the following command: `/usr/sbin/showmount -a`.
3. Add space management to your file system using the **dsmmigfs add** command.
4. To export the NFS file system again, issue the following command:
`/usr/etc/exportfs -a .`
5. Instruct all NFS clients to mount the exported NFS file system again.

Related reference:

“dsmmigfs add and update” on page 128

Removing or reconfiguring a space-managing node in a GPFS cluster

Reconfiguring a General Parallel File System (GPFS) node can change the host name or the GPFS node number. Removing or reconfiguring a node properly does not jeopardize failover and recall.

About this task

To remove or reconfigure a space-managing node in a GPFS cluster, follow these steps:

Procedure

1. Transfer the space-managed file systems to another node with the **dsmmigfs takeover** command.
2. Uninstall the space management client from the node that is removed or reconfigured.
3. Remove or reconfigure the GPFS node.

What to do next

If you want the node to participate in space management again, do the following steps:

1. Reinstall the space management client on the node.
2. Use the **dsmmigfs takeover** command to regain control of the file systems.

Related reference:

“dsmmigfs takeover” on page 146

Changing the mount point of a space-managed file system

If you change the mount point of a space-managed file system, you must rename the file space on the IBM Spectrum Protect server.

About this task

If you rename the file space before you mount the file system at the new mount point, you must rename the old file space to the new file space name.

If you changed the mount point and migrated some files before you rename the file space, you must merge the old file space and the new file space. If the new mount point is the same as an existing file space, you must also merge the old file space with the existing file space, even if you did not yet migrate any files at the new mount point.

Mounting a space-managed file system at a mount point for which a file space does not exist

If you mount a previously space-managed file system at a new mount point, you must rename the file space on the IBM Spectrum Protect server. If you rename the file space before you migrate any files at the new mount point, and if the new mount point does not match an existing file space, you do not have to merge file spaces.

Procedure

1. Unmount the file system. For GPFS file systems, unmount the file system on all GPFS nodes.
2. On the IBM Spectrum Protect server, rename the file space to the new name by using the **RENAME FILESPACE** administrative command. The new file space name must match the mount point.
3. For GPFS file systems, rename the GPFS mount point by using the **mmchfs** command. If the new mount point is `/hsmnew`, enter the following command:

```
mmchfs -T /hsm_new
```

For more information about the GPFS **mmchfs** command, go to the IBM Spectrum Scale product documentation.

4. Mount the file system at the new mount point. For GPFS file systems, mount the file system on the space-managed file-system owner node first. Then mount the file system on other GPFS nodes.

Related information:

 [RENAME FILESPACE administrative command](#)

Mounting a previously space-managed file system at a new mount point for which a file space already exists

If a file space exists for the new mount point of a space-managed file system, you must merge the old file space and the new file space.

Before you begin

Merging file spaces requires a secondary IBM Spectrum Protect server.

About this task

When you change the mount point of a space-managed file system, you must rename the file space on the IBM Spectrum Protect server. If you changed the mount point and migrated some files before you rename the file space, you must merge the old file space and the new file space. If the new mount point matches an existing file space, you must also merge the old file space with the existing file space, even if you did not yet migrate any files at the new mount point.

Procedure

1. Unmount the file system. For GPFS file systems, unmount the file system on all GPFS nodes.
2. Merge the old file space with the new file space on the IBM Spectrum Protect server.
 - a. Export the HSM data in the old file space to a secondary IBM Spectrum Protect server by using the **EXPORT NODE** administrative command. Specify `FILEData=SPacemanaged`.
 - b. On the secondary IBM Spectrum Protect server, rename the file space to the new name by using the **RENAME FILESPACE** administrative command. The new file space name must match the mount point and the existing file space name on the IBM Spectrum Protect primary server.
 - c. Import the HSM data from the IBM Spectrum Protect secondary server to the file space on the primary server by using the **IMPORT NODE** administrative command. Specify `FILEData=SPacemanaged` and `MERGEfilespaces=yes`.
3. For GPFS file systems, you must rename the GPFS mount point. If the new mount point is `/hsmnew`, enter the following command:

```
mmchfs -T /hsm_new
```

For more information about the GPFS **mmchfs** command, go to the IBM Spectrum Scale product documentation.

4. Mount the file system at the new mount point. For GPFS file systems, mount the file system on the space-managed file-system owner node first. Then mount the file system on other GPFS nodes.

Related information:

-  [RENAME FILESPACE administrative command](#)
-  [EXPORT NODE administrative command](#)
-  [IMPORT NODE administrative command](#)

Adding space management to workload partitions on AIX V6.1 and V7.1 operating systems

AIX V6.1 and V7.1 workload partitions (WPARs) act and look like stand-alone systems and provide an isolated environment for enterprise applications and data. A space management client installed in the global partition has access to all file system data across WPARs.

You can add space management only to WPAR file systems that are identified in the `/etc/filesystems` file.

To add space management support for WPARs, use the **dsmmigfs add** command. For example, to add space management to the file systems in two WPARs (`wpar1` and `wpar2`), use the following commands:

```
dsmmigfs add /wpars/wpar1/home
dsmmigfs add /wpars/wpar2/data
dsmmigfs add /home
dsmmigfs add /opt
```

When migrated files under `/wpars/wpar1/home` and `/wpars/wpar1/data` are accessed from their corresponding WPARs, the files are transparently recalled. They are transparently recalled if accessed from the global partition.

The space management client is supported only in the global partition.

Related information:

 [Backup-archive client processing options](#)

Configuration by using the command line

To update space management settings for your file system from the command line, use the **dsmmigfs update** command.

On GPFS file systems, you can use the GPFS policy engine to monitor space thresholds and search for migration candidates.

For information about configuring GPFS integration with the space management client, see Technote 7018848.

If you do not use the GPFS policy engine to control migration, specify the settings with the **dsmmigfs update** command.

Related reference:

“**dsmmigfs add** and **update**” on page 128

Settings that control the space usage of a file system

You can control the space usage of a file system by indicating how and when files are migrated and recalled.

You can control the following aspects of file migration and recall on a file system:

- The high and low thresholds for your file system that determine when threshold migration automatically starts and stops
- The total number of megabytes of data that you can migrate and premigrate from your file system to IBM Spectrum Protect storage
- The file size before it is migrated
- The size of the stub files that remain on your local file system when you migrate your files
- The order in which eligible files automatically migrate from your local file system
- The amount of free space the space management client maintains on your local file system
- The minimum size (in megabytes) for a file to qualify for partial file recall
- The minimum size (in megabytes) for a file to qualify for streaming recall mode
- The maximum number of files in the space-managed file system.

The information for each file system is stored in the `hsmfsconfig.xml` file, which is in the appropriate `.SpaceMan` directory. You can modify the settings in the

hsmfsconfig.xml file with the **dsmmigfs update** command. You can also use the **dsmmigfs query** command to display the settings in this file.

At any time after you add space management to your local file systems, you can update the settings, if necessary.

Copy several files into your migrated file system after you add space management and run the **dsmmigrate** command. If you are running open registration, the command prompts you for your node password and contact information the first time that you run it.

You can increase file transfer performance by allowing failover of automatic migration and recall capabilities to source nodes within a cluster environment.

Changes to space management settings take effect in the following manner:

- If you change the high and low thresholds or the premigration percentages for a file system, the new values take effect immediately.
- If you change the stub file sizes, the new values take effect only for files that are migrated after you change the settings.
- If you reduce the quota, and the data that you currently migrate exceeds the new quota, any additional files from your file system do not migrate. Sufficient files must be recalled during automatic or selective recall to drop the total number of megabytes of migrated and premigrated data below the new quota.
- If you change the maximum number of files and sufficient space is available, the current complete file index (CFI) is replaced by the newly sized CFI. Sufficient space is determined by the available free space plus the size of the current CFI. If the new value of the **maxfiles** parameter of the **dsmmigfs** command is 0, then the CFI is sized to the maximum required space in the file system. If the value of the **maxfiles** parameter is smaller than the actual number of files in the file system, or greater than the theoretical limit, then the CFI is sized to the theoretical limit for the file system.

On GPFS file systems, you can use the GPFS policy engine to monitor space thresholds and search for migration candidates.

For information about configuring GPFS integration with the space management client, see Technote 7018848.

Related concepts:

Chapter 3, “Configuring the space management client,” on page 19

Related reference:

“**dsmmigfs query**” on page 139

“**dsmmigfs add** and **update**” on page 128

“**dsmmigrate**” on page 149

Minimum migration file size

The space management client does not migrate a file unless the migration saves space on your local file system.

To be eligible for migration, a file must satisfy all of the following size criteria:

- The file logical size is larger than the value of the **stubsizes** option.
- The space that a file occupies is at least as great as the value of the **minmigfilesize** option.

The **minmigfilesize** option can be set for all file systems in the `dsm.sys` file. The **minmigfilesize** option can be set for a specific file system with the **dsmmigfs add** and **dsmmigfs update** commands. A valid setting for a specific file system overrides the global value that is specified in the `dsm.sys` file.

Related reference:

“**dsmmigfs add** and **update**” on page 128

“**minmigfilesize**” on page 107

Migration threshold percentages

The high and low threshold percentages for your file system affect when threshold migration starts and stops. A high threshold determines when threshold migration starts. A low threshold determines when file migration stops.

Specify a value of 0 through 100 percent. The default for a high threshold is 90 percent. The default for a low threshold is 80 percent. For example, if you allocate 10 GB for a file system, and you must maintain at least 1 GB of free space, set the high threshold to 90 percent. If space usage equals or exceeds 90 percent when the space management client checks space usage on your file system, files automatically begin migrating to IBM Spectrum Protect storage. The space management client migrates files beginning with the first file that is listed in the current migration candidates list for your file system.

The percentage that you specify for a low threshold must be the same as, or lower than, the percentage that you specify for a high threshold. For example, to stop migrating files when there are 2 GB of available free space on your file system, set the low threshold to 80 percent.

The realistic-minimum low-threshold percentage consists of the minimum file space usage (the percentage of the blocks used in the file system, if every file is migrated), the premigration percentage, and the percentage of file system space that is occupied by the `.SpaceMan` directory. For example, if the minimum file space is 55%, the premigration percentage is 10%, and the `.SpaceMan` directory occupies 10% of file system space, then the realistic-minimum low threshold is 75%. If the low threshold is set to 70%, then automatic migration tries to premigrate the 10% but it premigrates only 5%.

If the high threshold and the low threshold are the same, space usage must exceed the low threshold before threshold migration begins. When setting the realistic-minimum low-threshold option, remember that the `.SpaceMan` directory (which contains system-specific control files created by IBM Spectrum Protect) occupies some of the file system space. Files from this directory are not eligible for migration. You receive a `ANS9094W` message during automatic migration if the realistic-minimum low threshold does not fully include the percentage of space that is occupied by the `.SpaceMan` directory.

If there are no additional candidates in the migration candidates list after threshold migration starts, and if space usage drops below the high threshold that you set, threshold migration stops. The **dsmscoutd** daemon builds a new migration candidates list when candidates are available on your file system. Threshold migration starts again the next time your file system exceeds the high threshold.

On GPFS file systems you can use the GPFS policy engine to monitor space thresholds and search for migration candidates. If the GPFS policy engine controls automatic migration, ensure that the GPFS policy is sufficient to avoid an out-of-space condition for the space-managed file systems.

For information about configuring GPFS integration with the space management client, see Technote 7018848.

Related reference:

“**dsmmigfs add** and **update**” on page 128

Premigration percentage

The premigration percentage is the space that is occupied by premigrated files as a percentage of the total file system space. The **pmpercentage** option of the **dsmmigfs** command sets a target value for premigration.

The default value for the **pmpercentage** option is the difference between the percentages that you specify for the low and high thresholds. The default cannot be greater than the low threshold.

You can change the value of the **pmpercentage** option at any time.

When the premigration percentage is configured manually, the low threshold cannot be set to a value smaller than the premigration percentage.

If the percentage that you specify for the low threshold is the same as the percentage for the high threshold, the default premigration percentage is 0. The space management client does not premigrate any files after threshold or demand migration completes.

On GPFS file systems, you can use the GPFS policy engine to monitor space thresholds and search for migration candidates.

For information about configuring GPFS integration with the space management client, see Technote 7018848.

Related reference:

“**dsmmigfs add** and **update**” on page 128

Quotas

Quotas determine the maximum number of megabytes of data that you can migrate and premigrate from your file system to storage.

When files premigrate, they use space on both your local file system and in storage. When files migrate, stub files use some of the space on your local file system.

You can specify a quota value from 0 through 9999999999999999.

- If you set the quota to 0 for your file system, files do not migrate to storage. Set the quota for your file system to a value that is large enough to accommodate projected growth.
- If you set the quota to 9999999999999999, the amount of data you can migrate and premigrate is unlimited.
- The default is the number of MB that are assigned for your file system. For example, if 20 GB are assigned for your file system, the space management client migrates and premigrates your files from that file system until the total number of MB that migrate and premigrate equals 20 GB.

Check with your IBM Spectrum Protect administrator to determine whether there are any restrictions on the amount of data that you can migrate and premigrate to storage.

Related reference:

“**dsmmigfs add** and **update**” on page 128

Stub file size

When you migrate a file to IBM Spectrum Protect storage, a stub file is created on the local file system. A stub file contains information that is necessary to locate and recall a migrated file.

A stub file can contain leading bytes of data from the original file called *leader data*. You can read leader data without triggering a file recall.

More leader data requires more space on your local file system. More leader data can be useful if you frequently run programs that read only the information at the beginning of files.

For the space management client on GPFS file systems, you can specify 0 or a multiple of the file system block size. The default value is 0.

For all file system types, the maximum value for a stub file size is 1 GB.

Related reference:

“**dsmmigfs add** and **update**” on page 128

Minimum stream file size

Streaming recall mode allows for an asynchronous recall of migrated files. The recalled portion of the file can be accessed while the file is recalled.

Streaming recall mode is available in the following file system environments:

- AIX GPFS
- Linux GPFS

Streaming recall mode is valid for read-only operations on the file. The range of minimum stream file size value is 0 through 999999999. A value of 0 disables the asynchronous option and is the default.

Related reference:

“**dsmmigfs add** and **update**” on page 128

Minimum partial file size

In a GPFS environment, partial file recall recalls a portion of a migrated file. This avoids having to recall an entire, potentially large file, when only a small portion of the file is required by an application.

Related concepts:

“Recalling migrated files overview” on page 3

“File migration eligibility” on page 54

Maximum number of files

The maximum number of files determines the size of the complete file index (CFI). Changing the value of the maximum number of files changes the size of the CFI.

You can specify a maximum number of files value from 0 - 9999999999999999.

- If you change the maximum number of files and sufficient space is available, the current CFI is replaced by the newly sized CFI. Sufficient space is determined by the available free space plus the size of the current CFI.
- If the **MAXFiles** value is 0, the CFI is sized to the maximum theoretical limit for the file system. If the **MAXFiles** value is smaller than the actual number of files in the file system, or greater than the theoretical limit, then the CFI is sized to the theoretical limit for the file system.

If you use the GPFS policy engine to control automatic migration, this option is ignored.

For information about configuring GPFS integration with the space management client, see Technote 7018848.

Related concepts:

“Streaming recall mode” on page 73

Related reference:

“**dsmmigfs add** and **update**” on page 128

“**dsmmigfs query**” on page 139

Managing a file system with multiple IBM Spectrum Protect servers

Overview

A single General Parallel File System (GPFS) can be migrated to two or more IBM Spectrum Protect servers. Two or more IBM Spectrum Protect servers are required for file systems that contain more objects than can be managed by a single server. As the file system grows beyond the capacity of existing servers, you can add more servers to manage the file system.

A file system that is managed by multiple IBM Spectrum Protect servers must be GPFS version 3.4 or later and must be enabled for Data Management Application Programming Interface (DMAPI).

A file in a multiple-server environment is coupled with the IBM Spectrum Protect server to which it is initially migrated or backed up. The coupling between the file and the server is persistent over the life of both the file and the IBM Spectrum Protect server, and ensures data consistency. The file cannot be migrated or backed

up to another IBM Spectrum Protect server until the file is uncoupled. Automatic migration and transparent recall processes contact the server with which a file is coupled.

If a file has not been coupled with a server, you can specify the server when you do a selective migration of the file.

Automatic migration uses the GPFS policy engine to determine which IBM Spectrum Protect server to use for the initial migration. By default, the space management client distributes files to all servers in a round-robin fashion. You can customize the distribution scheme to suit your environment.

A single space management client can use only HSM daemons (**dsmonitord** and **dsmscoutd**) or GPFS policy to manage automatic migration. A space management client cannot use both HSM daemons and GPFS policy to manage automatic migration.

Related tasks:

“Restoring a file system in an environment that is managed by multiple IBM Spectrum Protect servers” on page 69

Enabling a file system to be managed by multiple IBM Spectrum Protect servers

Enable multiple-server support for each space management client in the GPFS node, then add IBM Spectrum Protect servers to manage the file system.

Before you begin

In a multiple-server environment, the IBM Spectrum Protect server to which a file is backed up must be the server to which a file is migrated. Before the upgrade, each file system on a GPFS cluster that is managed by a space management client must be managed by a single IBM Spectrum Protect server. That means that files on a file system must be backed up and migrated to the same IBM Spectrum Protect server.

A file system that is managed by multiple IBM Spectrum Protect servers must be GPFS version 3.4 or later and must be enabled for Data Management Application Programming Interface (DMAPI).

A single space management client can use only HSM daemons (**dsmonitord** and **dsmscoutd**) or GPFS policy to manage automatic migration. A space management client cannot use both HSM daemons and GPFS policy to manage automatic migration.

Multiple-server support requires that automatic migration is driven by the GPFS policy engine. Before you upgrade to multiple-server support, you must configure GPFS appropriately.

For information about configuring GPFS integration with the space management client, see Technote 7018848.

For information about using the GPFS backup command **mmbackup** with a file system that is space-managed by multiple IBM Spectrum Protect servers, go to the General Parallel File Systems product information and see **mmbackup** command: IBM Spectrum Protect requirements.

The following steps enable a file system to be managed by multiple servers.

Procedure

1. Set the **hsmdisableautomigdaemons=YES** option for each space management client in the GPFS cluster. Setting **hsmdisableautomigdaemons=YES** disables the automatic migration daemons **dsmonitor** and **dsmscoutd**. The **dsmonitor** and **dsmscoutd** daemons do not drive automatic migration in a multiple-server environment. If you do not disable the **dsmonitor** and **dsmscoutd** daemons, they continue to run and use system resources.
2. Set the **hsmmultiserver=YES** option for each space management client in the GPFS cluster.
3. Add to the list of servers the IBM Spectrum Protect server that currently manages the file system by issuing the following command: **dsmmigfs addmultiserver -server=server_name file_system_name**. You can determine the IBM Spectrum Protect server that currently manages the file system with the following command: **dsmmigfs query -detail**.
4. Run the `dsmMultiServerUpgrade.pl` script.
The script couples all files on a file system with the IBM Spectrum Protect server that manages the migration copies and backup versions. The script calls the **dsmreconcile** command to run a special reconciliation that couples the files with the server.
5. Optional: Add more servers to manage a file system. For each additional server, issue the following command: **dsmmigfs addmultiserver -server=server_name file_system_name**.

What to do next

To query status of the multiple-server environment, issue the following command: **dsmmigfs querymultiserver file_system_name**. You can add more IBM Spectrum Protect servers to the list of servers that manage a file system.

Related reference:

“**dsmmigfs addmultiserver, querymultiserver, and removemultiserver**” on page 134

“**hsmmultiserver**” on page 103

“**hsmdisableautomigdaemons**” on page 94

Adding space management to a file system in a multiple-server environment

You can add space management to a file system that is not currently space-managed. In a multiple-server environment you must first add space management, then add support for multiple IBM Spectrum Protect servers.

About this task

This task assumes that you enabled multiple-server support for another file system on the HSM node. This task also assumes that you add space management to a file system that is not currently space-managed.

Procedure

1. Add space management to the file system with the following command:
dsmmigfs add -server=server_name file_system_name options.

2. Add support for multiple servers with the following command: **dsmmigfs addmultiserver -server=server_name file_system_name**. You must specify the same IBM Spectrum Protect server in the **server** option in steps 1 on page 47 and 2.

Results

The space management client now manages space on the file system.

What to do next

To query status of the multiple-server environment issue the following command: **dsmmigfs querymultiserver file_system_name**. You can add more IBM Spectrum Protect servers to the list of servers that manage the file system.

Related reference:

“**dsmmigfs add** and **update**” on page 128

“**dsmmigfs addmultiserver**, **querymultiserver**, and **removemultiserver**” on page 134

Removing IBM Spectrum Protect servers from a multiple-server environment

You can remove one or more IBM Spectrum Protect servers from an environment that supports multiple servers.

Before you begin

Before you remove a server from managing a file system, recall all files from the IBM Spectrum Protect server to the local file system. Then, remove the server with the following steps:

Procedure

1. Remove the IBM Spectrum Protect server from the multiple server list by issuing the following command: **dsmmigfs removemultiserver -server=server_name file_system_name**. Run the command for all file systems that were managed by the server. After you run the command, files are not backed up or migrated to the server that was removed. You can still recall migrated copies and restore backup copies from the server that was removed.
2. Run the `dsmMultiServerRemove.pl` script.

The `dsmMultiServerRemove.pl` script does the following three tasks:

- a. Recalls all migrated files that are coupled with the IBM Spectrum Protect server that was removed to the specified file system.

The space management client determines how much space is required to recall all files to the file system. If there is not enough space, the space management client notifies you. To remove the server, make space available and run the script again.

- b. Reconciles the specified file system with the IBM Spectrum Protect server that was removed. The following reconciliation tasks are completed:
 - The space management client verifies that valid objects for each local stub file exist in the space management pool of the specified IBM Spectrum Protect server.
 - When orphans are identified, their names are recorded in the `.SpaceMan/orphan.stubs` file.

- The space management client expires all files in the space management storage pool on the IBM Spectrum Protect server.

Note: The script does not expire backup copies. If you want to also expire backup copies, you can modify the script by adding the backup-archive client command **dsmc expire**.

The **dsmMultiServerRemove.pl** script calls the `dsmreconcileGPFS.pl` script for the reconciliation.

- c. Uncouples all files that are coupled with the IBM Spectrum Protect server that was removed.

More than one recall and reconciliation process can run simultaneously. For example, you can simultaneously remove server `TSM_server` from several file systems with commands like the following commands:

```
dsmMultiServerRemove -server=TSM_server file_system_A
dsmMultiServerRemove -server=TSM_server file_system_B
dsmMultiServerRemove -server=TSM_server file_system_C
```

What to do next

To query status of the multiple-server environment, issue a command like the following command: **dsmmigfs querymultiserver** *file_system_name*.

After files are decoupled from the IBM Spectrum Protect server that was removed, they can be migrated and backed up to another server.

If at least one other IBM Spectrum Protect server manages the file system, the file system is still space-managed.

Related reference:

“**dsmmigfs addmultiserver**, **querymultiserver**, and **removemultiserver**” on page 134

“**dsmRemoveServer.pl**” on page 171

Limitations for multiple-server environments

Some features are not supported when a file system is managed by multiple IBM Spectrum Protect servers.

IBM Spectrum Protect server name encryption

It is not possible to encrypt multiple IBM Spectrum Protect server names on a single node.

IBM Spectrum Protect server node replication

All IBM Spectrum Protect client nodes in a General Parallel File System (GPFS) cluster have the same node name. The IBM Spectrum Protect server node replication feature requires that different client node names are used on the source and target servers. Therefore, you cannot use the node replication feature in a multiple-server environment.

IBM Spectrum Protect server LAN-free configuration

A LAN-free IBM Spectrum Protect server configuration is not supported in a multiple-server environment.

Replicating nodes in a multiple-server environment

If the primary IBM Spectrum Protect server for the space management client is unavailable, you can manually configure the space management client to connect to a secondary server. You can recall files from the secondary IBM Spectrum Protect server but cannot migrate files to the secondary server.

Before you begin

The primary IBM Spectrum Protect server for the space management client must be one that replicates client node data.

In a multiple-server environment, each IBM Spectrum Protect server includes storage pools for backup and migration. All pools used for backup and migration must be replicated to the same secondary server.

About this task

The IBM Spectrum Protect server that the client connects to during normal production processes is called the *primary server*. When the primary server is set up for node replication, the data for client nodes can be replicated to the *secondary server*.

The backup-archive client can automatically fail over to the secondary server when it is configured for failover.

The space management client, however, does not automatically fail over to the secondary server. You must manually edit the `dsm.sys` file to connect to the secondary server. Any secondary server information in the `replservername` stanza and `myreplicationserver` option is ignored by the space management client.

Procedure

1. Stop the backup and HSM processes on the GPFS cluster.
2. Edit the `dsm.sys` options file.
 - a. Change the value of the `servername` option from the primary server to the secondary server.
 - b. Do not change the name of the server stanza.
3. Start the HSM processes on the GPFS cluster.

Results

Functions that require write access to the secondary IBM Spectrum Protect server are not available for a space managed file system. You cannot back up, archive, or migrate files to the secondary server from a space managed file system. You can restore, retrieve, and recall from the secondary server.

With the backup-archive client, you cannot restore migrated files as stub files. Migrated files can only be restored as resident files.

What to do next

To switch back to the primary IBM Spectrum Protect server, repeat the steps, but change the value of the `servername` option from the target server to the source server.

Related information:

Deactivating space management

To deactivate space management on a single file system, use the **dsmmigfs** command with the **-deactivate** parameter. To deactivate space management on all space-managed file systems, use the **dsmmigfs** command with the **-globaldeactivate** parameter. Deactivate space management before you or your IBM Spectrum Protect administrator perform system maintenance to temporarily prevent migration, recall, or reconciliation processes from occurring. When you reactivate space management on your file systems, all space management services resume.

About this task

Deactivate space management before you or your IBM Spectrum Protect administrator perform system maintenance to temporarily prevent migration, recall, or reconciliation processes from occurring. When you reactivate space management on your file systems, all space management services resume.

Note:

- When you deactivate space management for a file system, the file system state becomes deactivated. Any migration, recall, or reconciliation process that currently is in progress completes before deactivation.
- You can access only resident and premigrated files on a deactivated file system.
- If your administrator exports migrated files from one server and imports them to another, update the `dsm.sys` file. The client node must contact the new server for space management services. The administrator can use a **lock node** command to prevent the client node from migrating or recalling files while importing and exporting to another server.

Related reference:

“dsmmigfs deactivate, reactivate, and remove” on page 137

“dsmmigfs globaldeactivate and globalreactivate” on page 138

Reactivating space management

If you deactivated space management from your file systems, you can reactivate space management by using the **dsmmigfs** command with the **-reactivate** parameter or the **-globalreactivate** parameter.

Related reference:

“dsmmigfs deactivate, reactivate, and remove” on page 137

“dsmmigfs globaldeactivate and globalreactivate” on page 138

Removing space management

You can remove space management from a file system by using the `thedsmmigs` command with the `-remove` parameter.

Before you begin

Space management must be active on your file system to completely remove it. You can remove space management with a single command. Before you remove space management from your file system, ensure that you have enough space on your file system to recall all migrated files. Ensure also that all activity on your file system is stopped, and that the file system is not being accessed.

About this task

When you remove space management from your file system, the space management client does the following tasks:

- Runs reconciliation for your file system. If any orphaned stub files are located, a notification is posted. Check the `orphan.stubs` file in the `.SpaceMan` directory for a list of orphaned stub files. To remove space management, first resolve all orphaned stub files, and then try again.
- Determines how much space is required to recall all migrated files. If there is not enough space, the space management client notifies you. To remove space management, make space available and try again.
- Recalls migrated files to the file system in the most efficient recall order that is based on where they are stored.
- Notifies the server to delete all migrated files from storage.
- Deletes the `.SpaceMan` directory from your file system.
- Updates information for your native file system.

Related concepts:

“Orphaned stub files” on page 81

Related reference:

“`dsmmigs deactivate`, `reactivate`, and `remove`” on page 137

Chapter 5. Migrating files

There are several types of file migration. You can configure the space management client to migrate files when a files system is low on free space, or you can selectively migrate files at any time. A file must meet several criteria to be eligible for selective or automatic migration.

Related information:

 [Server: Data protection strategies with IBM Spectrum Protect](#)

Migration types

The space management client provides automatic and selective migration. Automatic migration is triggered by space-usage thresholds (threshold migration) or an out-of-space condition (demand migration).

Automatic migration monitors space usage and automatically migrates eligible files according to the options and settings that you select. The space management client provides two types of automatic migration: threshold migration and demand migration.

Threshold migration

Threshold migration maintains a specific level of free space on your local file system. When space usage reaches the high threshold that you set for your file system, eligible files are migrated to IBM Spectrum Protect server storage automatically. When space usage drops to the low threshold that you set for your file system, file migration stops.

On GPFS file systems, you can use the GPFS policy engine to monitor space thresholds and search for migration candidates.

For information about configuring GPFS integration with the space management client, see Technote 7018848.

Demand migration

Demand migration responds to an out-of-space condition on your local file system. Demand migration starts automatically if your file system runs out of space. The process or event that caused the out-of-space condition continues. On GPFS file systems, the process receives an out-of-space (ENOSPC) return code and stops.

Selective migration moves specific files from your local file system to IBM Spectrum Protect server storage. You specify the files to migrate. Migration is not triggered automatically by file-system free space.

File migration eligibility

A file must meet several criteria to be eligible for migration.

A file is eligible for automatic migration, selective migration, or premigration when it meets the following criteria:

- The file is a regular file. Character special files, block special files, FIFO special files (named pipe files), or directories are not migrated.
- The file is a resident or premigrated file on a file system for which space management is active.
- The file is not excluded from space management in your include-exclude options file.
- To be eligible for migration, a file must satisfy all of the following size criteria:
 - The file logical size is larger than the value of the **stubsiz**e option.
 - The space that a file occupies is at least as great as the value of the **minmigfilesize** option.

Restriction: The size criteria do not apply for premigration.

- The file meets management class criteria.
- A file in a multiple-server environment that is coupled with a server can be migrated to only that server.
- The file is more than 2 minutes old. You can migrate files that are less than 2 minutes old if you set **hsmenableimmediatemigrate** = yes.

You can migrate a file in a set of hard linked files. The file must not be excluded from space management and must be assigned to a management class that allows automatic or selective migration.

Note: Any application that touches a file causes the last access date of the file to change. The last access date is one of the factors that determines when a file becomes eligible for migration.

The IBM Spectrum Protect backup-archive client uses the **preserve**lastaccessdate option to determine whether to reset the last access date of any specified files after a backup or archive operation. By default, the backup-archive client does not reset the last access dates to their original values after the backup or archive operation.

Related concepts:

Chapter 1, “Overview of the space management client,” on page 1

“Migrating files overview” on page 2

Related reference:

“**dsmdf**” on page 123

“**dsmdu**” on page 124

“**hsmenableimmediatemigrate**” on page 95

File premigration

For faster migration, the space management client prepares files for automatic migration by using a process that is called *premigration*.

Premigrated files are copied to IBM Spectrum Protect server storage while the original files remain on your local file system. The next time that you need free space on your local file system, the space management client just changes the status of premigrated files to migrated files. No additional time is required to copy the files to storage.

The space management client verifies that files did not change since they became premigrated. When your premigrated files migrate to IBM Spectrum Protect server storage, stub files replace them on your local system.

The space management client premigrates files each time it completes automatic migration if the following conditions are true:

- The file system contains extra files that are eligible for automatic migration.
- The premigration percentage that you set for your file system is not reached or exceeded.

For example, if the high threshold is 90 percent and the low threshold is 80 percent, the premigration target default value is 10 percent. When space usage reaches the high threshold, files are migrated until space usage is 80 percent of the file system total space. Next, files are premigrated until at least 10 percent of the occupied space on your file system contains premigrated files. The next time that threshold migration is required, the space management client replaces the premigrated files with stub files. Space usage is quickly reduced to the low threshold. No files must be copied to IBM Spectrum Protect server storage.

Related reference:

“**hsmgroupedmigrate**” on page 97

“**dsmmigrate**” on page 149

Automatic file migration

To ensure that free space is available on your local file systems, the space management client monitors space usage and automatically migrates files whenever it is necessary. Files are prioritized for automatic migration based on the age and size settings. The **dsmscoutd** daemon searches these files in cycles and upon request from automatic migration.

The space management client provides two types of automatic migration: threshold and demand.

Threshold migration maintains a specific level of free space on your local file system. The space monitor daemon checks space usage on your local file systems at intervals that you specify. When space usage reaches the high threshold that you set for a file system, migration automatically sends eligible files to IBM Spectrum Protect server storage. When space usage reaches the low threshold that you set for a file system, migration stops. For example, assume that you set the high threshold for your file system to 80 percent and the low threshold to 70 percent. Files begin migrating to IBM Spectrum Protect server storage when there is less than 20 percent of available space on your local file system. Files stop migrating when there is more than 30 percent of available space on your local file system.

More files are premigrated as specified with the **pmpercentage** parameter of the **dsmmigfs** command.

Demand migration responds to an out-of-space condition on your local file system. The space monitor daemon checks for an out-of-space condition every 10 seconds. Threshold migration starts automatically if the used capacity of your file system exceeds a certain limit. The default for this limit is 90 percent. For example, assume that you attempt to copy a large file into your file system, and there is not enough available space for the file. Eligible files begin migrating automatically from your local file system to IBM Spectrum Protect server storage. As space becomes available, the process continues to copy the large file to your file system.

Demand migration is not supported on GPFS file systems. Copying a large file to a GPFS file system with insufficient space yields an out-of-space error message.

Hidden directories and files are included in automatic migration. Hidden objects can be excluded from automatic migration by adding the hidden objects to the exclude list in the `dsm.sys` file.

There is a potential impact on applications that depend on a timely response to write requests. The delay time depends on how fast objects are migrated from the file system to create free space and on the configuration of the space management client. For instance, if the HSM client is configured to require a backup before migration, the migration process can be delayed until the objects are backed up. To avoid a long delay, you must have enough eligible migration candidates.

On GPFS file systems, you can use the GPFS policy engine to monitor space thresholds and search for migration candidates. In this case, the space monitor and scout daemons are disabled.

For information about configuring GPFS integration with the space management client, see Technote 7018848.

Related reference:

"dsmmigfs add and update" on page 128

Candidate selection for automatic migration

Files in a file system, which become eligible for automatic migration, are considered candidates for automatic migration. The **dsmscoutd** daemon finds the best candidates for automatic migration. It also prioritizes the files based on the file age, the number of days since a file was last accessed, and the file size.

For a file to be eligible for automatic migration, it must meet these requirements:

- Be in a file system to which space management was added
- Meet all management class requirements for eligibility
- Meet the minimum required size for migration
- Be included for space management services

When a file system is added to space management, the **dsmscoutd** daemon creates a complete file index (CFI) in the `.SpaceMan/metadata` directory for the managed file system. The CFI requires space equivalent to 1% to 3% of the managed file system. If the CFI is not created, view the `dsmerror.log` file for more information and use the **dsmscoutd scanplan** command to obtain additional information.

When a file system is removed from space management, the CFI is removed.

The CFI is updated by various processes:

- The daemon periodically scans the managed file systems and updates the CFI with the latest information about every file. The information is used to generate a list of files that is used for threshold and demand migration. Before migration, each file is checked to determine whether it is still eligible for migration.
- The **dsrecall** and **dsmigrate** commands update the CFI. When a file is restored as a migrated file, the file information is stored in the CFI.

The following commands and options can modify or monitor automatic migration:

- The **minmigfilesize** option in the `dsm.sys` file. Specify the minimum file size for a file to be eligible for automatic migration.
- The **dsmscoutd scanplan** command. Monitor the next scan time or the remaining time before the next scan for one or more managed file systems.
- The **maxcandprocs** option in the `dsm.sys` file. Change the number of scans that can run in parallel in the daemon.
- The **maxcandidates** parameter of the **dsmigfs** command. Improve the performance of automatic migration by lowering the value to 100.

Note: Do not increase the value of **maxcandidates** higher than 1000. A higher value slows down the automatic migration. For best results on your system, select a value from 10 - 500. The optimal number to use depends on the performance of the file system and operating system.

On GPFS file systems, you can use the GPFS policy engine to monitor space thresholds and search for migration candidates. In this case, the space monitor and scout daemons are disabled.

For information about configuring GPFS integration with the space management client, see Technote 7018848.

Related concepts:

“The scout daemon” on page 84

Related reference:

“**minmigfilesize**” on page 107

“**maxcandprocs**” on page 104

Manually starting threshold migration

Use the **dsmautomig** command to reduce space usage to the low threshold on your file system before it reaches the high threshold. If space usage exceeds the low threshold when you start threshold migration manually, eligible files migrate until space usage drops to the low threshold.

About this task

For one or more storage pools that are configured in a file system, the low and high thresholds defined for a file system also apply to each storage pool in that file system.

Each storage pool is monitored and managed separately. Running automatic migration on one pool until it reaches the low threshold does not result in a low threshold for the entire file system. To reach low threshold for the entire file system, issue the **dsmautomig** command for the entire file system (without a storage pool argument). Low threshold is enforced for each storage pool in that file system.

Files are automatically premigrated after threshold migration completes. There must be files that are eligible for premigration and the number of premigrated files must not exceed the premigration percentage. The status of the premigration process is displayed in the Threshold Migration Status window.

Procedure

1. To start threshold migration from the command line, use the **dsmautomig** command. For example, to start threshold migration for the /home file system, issue the following command:

```
dsmautomig /home
```
2. To display information about your migrated files, use the **-detail** parameter with the **dsmautomig** command.

Related reference:

“**dsmautomig**” on page 121

Selective file migration

Use selective migration to move specific files from your local file systems to storage.

Automatic migration occurs as a response to the lack of free space on a file system. Selective migration has no dependence on file system free space. You can start selective migration at any time.

Migrating selectively by using the **dsmmigrate** command

Use selective migration (the **dsmmigrate** command) to move specific files from your local file systems to IBM Spectrum Protect server storage. For example, if you are not using some files for an extended time, you can migrate them to storage to free extra space on your local file system.

About this task

Note: On large file systems, selective migration can take a while to complete. The process can be quicker if you migrate only premigrated files by using the **stubpremigrated** option.

The space management client migrates files that are eligible for selective migration according to the settings and options you define. When you migrate a file selectively, the access time for the file does not change. Unlike automatic migration, the number of days since you last accessed a file has no effect on whether your file is eligible for selective migration.

If you plan to migrate many small and medium size files (up to 100 MB) directly to tape, set the **hsmgroupedmigrate** option to **YES**. With the **HSMGROUPedmigrate** option, you can activate the HSM transaction grouping to improve migration performance for bulk operations that move small and medium size files directly to tape.

An eligible file must meet the following management class requirements:

- The management class that is assigned to the file enables selective migration.
- A current backup version of the file must exist on your migration server if the management class requires one.

For command syntax, options, and examples, see the **dsmmigrate** command.

Related tasks:

“Migrating selectively by using the **dsmmigrate** command” on page 58

Related reference:

“**dsmmigrate**” on page 149

“**hsmgroupedmigrate**” on page 97

Migration of a single file system to two or more IBM Spectrum Protect servers

A single General Parallel File System (GPFS) can be migrated to two or more IBM Spectrum Protect servers. Two or more IBM Spectrum Protect servers are required for file systems that contain more objects than can be managed by a single server. As the file system grows beyond the capacity of existing servers, you can add more servers to manage the file system.

A file system that is managed by multiple IBM Spectrum Protect servers must be GPFS version 3.4 or later and must be enabled for Data Management Application Programming Interface (DMAPI).

Related concepts:

“Managing a file system with multiple IBM Spectrum Protect servers” on page 45

Chapter 6. Backup and restore on space managed file systems

Coordinate backup and migration to protect the file system data.

When you back up a file with the backup-archive client, a copy of the file is created on the IBM Spectrum Protect server. The original file remains in your local file system. To obtain a backed up file from IBM Spectrum Protect server storage, for example in case the file is accidentally deleted from the local file system, you restore the file. In contrast, when you archive a file to IBM Spectrum Protect server storage, that file is removed from your local file system, and if needed, you retrieve it from IBM Spectrum Protect storage.

When you migrate a file, you move the file to IBM Spectrum Protect server storage and replace it with a stub file on your local file system. You can then use that stub file to recall the full file from its migration location.

Regularly use the IBM Spectrum Protect backup-archive client to guard against loss or corruption of your data, regardless of whether the files are resident, migrated, or premigrated.

You can back up and migrate your files to the same IBM Spectrum Protect server or to different IBM Spectrum Protect servers. If you back up and migrate files to the same server, the HSM client can verify that current backup versions of your files exist before you migrate them. For this purpose, the same server stanza for backup and migration must be used. For example, if you are using the **defaultserver** and **migrateserver** options, they must both point to the same server stanza within the `dsm.sys` file. You cannot point to different server stanzas, even if they are pointing to the same IBM Spectrum Protect server.

To restore stub files rather than backup versions of your files, use the backup-archive client **restore** command with the **restoremigstate** option. To restore the stubs of space-managed files with the backup-archive client, the **ds mreca11d** daemon must be running. Your migrated and premigrated files remain intact on the IBM Spectrum Protect server, and you restore only the stub files on your local system. You cannot use the backup-archive client to restore stub files for your migrated files, if the files were backed up before the files were migrated. In this case, use the space management client **ds mmigundelete** command to re-create stub files for any migrated or premigrated files that are lost.

If you move files after they are migrated and then back up, the migration destination of a migrated file can be different from a premigrated file. If you move a migrated file within the same file system and then back up, the file is moved on the IBM Spectrum Protect server to the new location. If you move a premigrated file within the same file system and then back up, the file is not automatically moved on the IBM Spectrum Protect server. You can move the migration location of the premigrated files when you reconcile the file system and check for orphan files. You can use either of the following reconcile commands:

1. Two-way orphan check reconciliation: `ds mreconcileGPFS.pl /file_system`
2. Orphan check reconciliation: `ds mreconcile -o /file_system`

If you back up and migrate data to tape volumes in the same library, ensure that there are always some tape drives available for space management. You can limit the number of tape drives that can be used simultaneously by backup and archive operations. Specify a number for the `mountlimit` that is less than the total number of drives available in the library. Using disk storage as your primary storage pool for space management might, depending on the average size of your files, result in better performance than using tape storage pools.

For information about the `mountlimit` option of the **define devclass** command, see the IBM Spectrum Protect Administrator's Reference for your operating system.

The space management client cannot verify that current backup versions of your files exist before a migration job in two cases:

- You back up files to one IBM Spectrum Protect server and migrate them to a different server.
- you use different server stanzas for backup and migration

In these cases, use the backup-archive client to restore complete backup versions.

Related concepts:

Chapter 1, "Overview of the space management client," on page 1

Related information:

 [Archive and retrieve data with backup-archive clients](#)

Archive and retrieve of migrated files by using the backup-archive client

You can archive your migrated and premigrated files at any time and retrieve them to your local file systems when you need them.

Related information:

 [Archive and retrieve data with backup-archive clients](#)

Migration status after a migrated file is archived

A migrated or premigrated file can be archived. The migration status depends on the storage location.

If you archive a migrated file to the same IBM Spectrum Protect server to which it was migrated, the file is not recalled to the local file system. The file is copied on the IBM Spectrum Protect server from the migration destination to the archive destination. If you archive a migrated file to another IBM Spectrum Protect server, the file is recalled before it is archived. The recalled file remains in premigrated state until the file is migrated again or the file becomes a resident file.

If you archive a premigrated file, the file is sent from your local file system to IBM Spectrum Protect storage.

If you erase a stub file from your local file system, the migrated copy of the file remains in IBM Spectrum Protect storage until the file expires. The file expires after the number of days that you specify on the **migfileexpiration** option in your `dsm.sys` file.

Migration status after a migrated file is archived and retrieved

A migrated or premigrated file can be archived and then retrieved. The migration status of the retrieved file is determined by the **restoremigstate** option in the `dsm.opt` file.

To restore archived files to resident status and remove the file from IBM Spectrum Protect storage after expiration, set the **restoremigstate** option to `NO`. The file is restored to the local file system in resident status.

To retrieve a file to your local file system and maintain a copy of the migrated file in storage, set the **restoremigstate** option to `YES`. The file is in premigrated status.

Note: Files with ACLs are restored to a resident state, even when you specify `YES` on the **restoremigstate** option.

Back up before migration

You must back up your migrated files to guard against data loss. Use a management class to specify whether your files are backed up before you migrate them.

If you back up and migrate files to the same server

Before a file is migrated, the space management client can check if a backup copy exists. Some management classes, including the default management class, require a backup copy before a file can be migrated. With the default management class, if a backup copy does not exist, the file is not migrated.

If you back up files to one server and migrate files to a different server

The space management client cannot verify that current backup versions of your files exist before you migrate files. Any management class that you assign to files must specify that current backup versions are not required before migration. Otherwise, you cannot migrate your files.

To back up your files after you migrate them, assign a management class to your files that does not include the requirement for an existing backup version. If you back up files to the same server to which you migrated them, files are copied from the migration destination to the backup destination. Files are not recalled to your local file system.

Backing up migrated files

Use the IBM Spectrum Protect backup-archive client to back up or archive migrated files. The **skipmigrated** option must not be set to `yes`.

Before you begin

The **skipmigrated** option controls whether migrated files are included in a backup or archive operation. If `skipmigrated=yes`, migrated files are excluded from the backup or archive operation. If `skipmigrated=no` (the default), the backup-archive client can process stub files during some operations.

Set the **skipmigrated** option in the backup-archive client options file `dsm.opt`. Use the preferences editor or directly edit the backup-archive `dsm.opt` options file. You can also specify the **skipmigrated** option when you start a command on the command line.

About this task

- If `skipmigrated=no` (the default), the type of backup or archive operation determines whether migrated files are backed up.

Incremental backup or image backup

Only an incremental backup or image backup can back up a stub. The object that is backed up depends on whether the IBM Spectrum Protect server contains a current backup copy of the complete file.

If the IBM Spectrum Protect server contains a current backup copy of the complete file:

An incremental or image backup backs up the stub.

If the IBM Spectrum Protect server does not contain a current backup copy of the complete file:

An incremental or image backup backs up the complete file.

Incremental-by-date backup

An incremental-by-date backup does not back up the stub or the complete file.

Selective backup or archive

Selective backup or archive does not back up a stub. The complete file is backed up regardless of whether a current backup copy exists on the IBM Spectrum Protect server.

The backup-archive client ensures that whenever a stub is backed up, there is a copy of the migrated file in the backup pool. If a migrated file was not backed up before migration, the migrated file is temporarily retrieved and is backed up. After the backup, the temporary file is removed by the backup-archive client. You can control the location to which the backup-archive client retrieves the temporary file by using the **stagingdirectory** option of the backup-archive client.

Tip: If you have many files that have not been backed up before migration, and the **skipmigrated** option has the default value `no`, all of those files are retrieved when they are backed up. The number of files that are retrieved during a backup can be great in these situations:

- You have many stubs that were backed up with backup-archive client version 5.4 and earlier versions. These files are temporarily retrieved during backups with later versions of the backup-archive client.
- You renamed or changed the security settings of stubs or directories containing stubs.
- You changed the backup policies for a volume by including for backup many files that were not previously included.

If the backup-archive client cannot create a backup copy of the migrated file, the backup-archive client does not back up the stub file. (The backup of the migrated file can be unsuccessful if the stub is an orphan.)

Restriction: You must not use adaptive subfile backup and HSM. You must back up only the entire migrated file. If you use adaptive subfile backup on migrated files, you might not be able to restore migrated files correctly. The backup-archive client does not report any errors or warnings when you do a subfile backup of a migrated file.

Procedure

Use the backup-archive client **archive**, **backup image**, **incremental**, or **selective** command.

For more information about these commands, see the IBM Spectrum Protect backup-archive client documentation:

- Archive command
- Backup image command
- Incremental command
- Selective command

Related reference:

“**skipmigrated**” on page 110

Restoring migrated files

Use the backup-archive client to restore a stub file or a complete file after the file was migrated. The **restoremigstate** option determines whether a stub or a complete file is restored.

Before you begin

A file must be backed up and migrated to the same IBM Spectrum Protect server.

If you set the **restoremigstate** option to **NO**, you restore a complete file to the local file system. The file has resident status. The migrated copy of the file is removed from IBM Spectrum Protect storage when the file expires.

If you set the **restoremigstate** option to **YES**, you restore a stub file to the local file system.

Note: Files with ACLs are restored to a resident state, even when you set **restoremigstate** to **YES**.

About this task

On a space-managed file system, you can restore a stub file only during standard-query restore processing. During no-query restore processing to a space-managed file system, you cannot restore stub files.

You can use the space management client **dsmmigundelete** command to restore stub files.

Procedure

Use the backup-archive client **restore** command.

For more information about the IBM Spectrum Protect backup-archive client **restore** command, see the IBM Spectrum Protect backup-archive client documentation: Restore command.

Related concepts:

“Re-create stub files by using the **dsmmigundelete** command” on page 66

Related reference:

“**restoremigstate**” on page 109

“**migfileexpiration**” on page 106

Re-create stub files by using the `dsmmigundelete` command

The `dsmmigundelete` command uses the migrated files on the IBM Spectrum Protect server to re-create stub files.

If your stub files are erased or corrupted, you can use the `dsmmigundelete` command to re-create stub files for all eligible migrated files.

Note:

- You cannot use the `dsmmigundelete` command to re-create stub files for individual files or specific groups of files.
- You can create stub files for any premigrated files for which an original file does not exist on your local file system.
- If backup versions of your stub files are available on the IBM Spectrum Protect server, use the backup-archive client to restore your stub files.
- The `dsmmigundelete` command does not support hard linked files. To re-create a stub file for a hard linked file, all files that are hard linked together must be deleted from your local file system. When one file in a set of hard linked files is migrated, all of the hard linked files in the set become stub files. When the `dsmmigundelete` command re-creates a stub file for a hard linked file, the stub file has the same name as the file that was originally migrated. Stub files are not re-created for any other files that were previously in the hard linked set of files.

The `dsmmigundelete` command has the following limitations:

- The `dsmmigundelete` command creates a stub file that contains the necessary information to recall the corresponding file from storage. The stub file does not contain any leading bytes of data from the file.
- The recall mode that you previously set for a migrated file is not stored in a re-created stub file. The recall mode for the file is set to normal.
- The space management client does not create a stub file if a directory path does not exist in your local file system for a migrated file.
- The space management client creates a stub file with the name of the file at the time it was migrated. If you rename a file after it is migrated, the file name is not updated on the server.
- If you have more than one migrated file in storage with the same name that is marked for expiration, the space management client creates a stub file again for the file with the most recent modification time (mtime).
- The space management client creates a stub file only if the modification time for the migrated file is newer than file on your local file system. If a file is renamed after it is migrated, the space management client creates a stub file regardless of the modification time.
- If you move files after they are migrated or premigrated and then the files are backed up, you can use the `dsmmigundelete` command to re-create stub files. The stub files are created at the location to which the files were moved.

The `dsmmigundelete` command re-creates a stub file for a migrated file or creates a stub file for a premigrated file even in the following situations:

- The file was never backed up.
- The migrated or premigrated file is on a different server other than the server that contains the backup copies of the file.
- The file was migrated or premigrated after the last incremental backup.

Issue the **dsmmigundelete** command without the **expiring** option if reconciliation was not run since the files were deleted. The space management client completes the following action for the file system that you specify:

Table 11. dsmmigundelete command actions without the expiring option

File	Description
Migrated files	Re-creates a stub file for a migrated file if a corresponding stub file does not exist on your local file system, and the migrated file was not marked for expiration. The migrated file is unexpired at the server after the dsmmigundelete command is finished.
Premigrated files	Creates a stub file for a premigrated file if a corresponding original file does not exist on your local file system, and the premigrated file was not marked for expiration. The premigrated file is unexpired at the server after the dsmmigundelete command is finished.

If you issue the **dsmmigundelete** command with the **expiring** option, the space management client completes the following action for the file system that you specify:

Table 12. dsmmigundelete command actions with the expiring option

File	Description
Migrated files	Re-creates a stub file for a migrated file if a corresponding stub file does not exist on your local file system. The stub is created whether the migrated file was marked for expiration or not.
Premigrated files	Creates a stub file for a premigrated file if a corresponding original file does not exist on your local file system. The stub is created whether the premigrated file was marked for expiration or not.

Related reference:

“**dsmmigundelete**” on page 153

Restoring space managed file systems

When you restore an entire file system, you can restore to the state as of the last incremental backup. Restoring stubs saves disk space and is faster than restoring the complete version of migrated files.

About this task

When you restore space managed files systems, if you restore many stubs with complete files, the file system can run out of space. If the file system runs out of space, the space management client must migrate files to make room for more restored files. The restore process can be slowed. To reduce disk space and complete the process quickly, choose one of the following tasks.

Related tasks:

“Restoring a file system that is backed up and migrated to the same IBM Spectrum Protect server” on page 68

“Restoring a file system after backup and migration operations used different IBM Spectrum Protect servers” on page 68

Restoring a file system that is backed up and migrated to the same IBM Spectrum Protect server

If you back up and migrate files to the same IBM Spectrum Protect server, you can restore the file system with resident files and stubs by using a single command.

About this task

Follow these steps to restore your file system in the exact order that is presented. Otherwise, you might not obtain the results that you want.

Procedure

1. Install the IBM Spectrum Protect backup-archive client, API, and space management client.
2. Configure the backup-archive client to connect to the IBM Spectrum Protect server where you backed up files from the file system. Configure the space management client to connect to the IBM Spectrum Protect server.
3. Follow your operating system instructions to establish the file system again.
4. Mount the file system.
5. Add space management to your file system.
6. Enter the backup-archive client **dsmc restore** command and set the **restoremigstate** option to YES (the default). For example, to restore the /home file system, enter the following command:

```
dsmc restore -restoremigstate=yes -sub=yes "/home/*"
```

For more information about the IBM Spectrum Protect backup-archive client **dsmc restore** command, see Restore command.

Results

The backup-archive client restores backup versions of resident files and restores stub files for migrated and premigrated files.

Related tasks:

“Adding space management to file systems” on page 36

Related reference:

“**dsmmigfs add** and **update**” on page 128

Related information:

 Restoring data

Restoring a file system after backup and migration operations used different IBM Spectrum Protect servers

You can restore a file system after you backed up files to one IBM Spectrum Protect server and migrated files to another server.

Procedure

1. Install the IBM Spectrum Protect backup-archive client, API, and space management client.
2. Configure the backup-archive client to connect to the IBM Spectrum Protect server where you backed up files from the file system. Configure the space management client to connect to the IBM Spectrum Protect server where you migrated files.
3. Follow your operating system instructions to establish the file system again.

4. Mount the file system.
5. Add space management to your file system.
6. Enter the backup-archive client **dsmc restore** command with the **dironly** option to restore the directory structure of your file system.

When you use the **dironly** option with the **dsmc restore** command, only directories are restored for your file system. For each directory, attributes such as access permissions or, on an AIX workstation an access control list, are restored. For example, to restore the directory structure for the /home file system, enter the following command:

```
dsmc restore -dironly -sub=yes "/home/*"
```

7. Enter the space management client **dsmmigundelete** command to re-create stub files for migrated files and to create stub files for premigrated files. For example, to restore stub files for all migrated and premigrated files in the /home file system that were not marked for expiration on the server, enter the following command:

```
dsmmigundelete /home
```

Note: The space management client restores hard links during the restore process if the hard links were backed up.

8. Enter the backup-archive client **dsmc restore** command with the **replace** option set to **NO** to restore backup versions of previously resident files. For example, to restore backup versions of all the remaining files in the /home file system, enter the following command:

```
dsmc restore -replace=no -sub=yes "/home/*"
```

Results

Files that were migrated or premigrated are restored as stub files from the IBM Spectrum Protect migration server. Files that were not migrated or premigrated are restored as complete files from the IBM Spectrum Protect backup server.

Related concepts:

“Re-create stub files by using the **dsmmigundelete** command” on page 66

Related tasks:

“Adding space management to file systems” on page 36

Related reference:

“**dsmmigfs add** and **update**” on page 128

“**dsmmigundelete**” on page 153

Related information:

 Restoring data

Restoring a file system in an environment that is managed by multiple IBM Spectrum Protect servers

You can recover files that were migrated and backed up from a single file system to multiple IBM Spectrum Protect servers.

About this task

This task assumes that you migrated and backed up from a single file system to multiple IBM Spectrum Protect servers.

Procedure

1. Determine the IBM Spectrum Protect servers that manage the file system. Issue the following command: **dsmmigfs querymultiserver** *file_system*.
2. Restore the directory structure. For each IBM Spectrum Protect server that manages the file system, issue the following command: **dsmc restore -dironly -latest -server=migration_server** *file_system_name*.
3. Restore the stub files. For each IBM Spectrum Protect server that manages the file system, issue the following command: **dsmmigundelete -server=migration_server** *file_system*. The **dsmmigundelete** command creates stub files and couples the stub files with the IBM Spectrum Protect server.
4. Restore files that were not migrated or premigrated. For each IBM Spectrum Protect server that manages the file system, issue the following command: **dsmc restore -ifnewer -server=migration_server** *file_system*. The **ifnewer** option ensures that the latest version of a file is recovered.

Related concepts:

“Managing a file system with multiple IBM Spectrum Protect servers” on page 45

Related reference:

“**dsmmigfs addmultiserver**, **querymultiserver**, and **removemultiserver**” on page 134

“**dsmmigundelete**” on page 153

Chapter 7. Recalling migrated files

A migrated file is recalled to the file system automatically when the file is accessed. You can also recall selected files without accessing the files.

A migrated file is not recalled when you access the file and all of the following conditions are true:

- The file is not modified by the operation.
- The stub contains all of the information that is required for the operation.

Related concepts:

Chapter 1, “Overview of the space management client,” on page 1

“Recalling migrated files overview” on page 3

Related reference:

“`dsmattr`” on page 118

Transparent recall processing

Transparent recall processing automatically returns a migrated file to its originating local file system when you access it. After the file is recalled, the space management client leaves the copy of the file in IBM Spectrum Protect server storage. The file becomes a premigrated file because an identical copy exists both on your local file system and in IBM Spectrum Protect server storage. If you do not modify the file, it remains premigrated until it again becomes eligible for migration.

A transparent recall process waits for a tape drive to become available. If you back up and migrate data to tape volumes in the same library, make sure that there are always some tape drives available for space management. You can achieve tape drive availability by limiting the number of tape drives that can be used simultaneously by backup and archive operations. Specify a number for the `mountlimit` that is less than the total number of drives available in the library. Using disk storage as your primary storage pool for space management might, depending on the average size of your files, result in a better performance than using tape storage pools.

For information about the `mountlimit` option of the **define devclass** command, see the IBM Spectrum Protect Administrator's Reference for your operating system.

If you modify a recalled file, it becomes a resident file. The next time your file system is reconciled, the space monitor daemon marks the stored copy for expiration.

Selective recall processing

Use selective recall processing if you want to return specific migrated files to your local file system. When you selectively recall a group of files, they are recalled in the most efficient, time-saving order based on where they are stored.

For example, if some of your files are stored on a disk storage device and some of your files are stored on a tape storage device, the space management client recalls all of your files that are stored on the disk storage device first. It next recalls all of your files that are stored on the tape storage device. The access time (atime) changes to the current time when you selectively recall a migrated file.

To selectively recall files, use the **dsmrecall** command. For example the following command recalls a file named proj1rpt to the /home/proja directory:

```
dsmrecall /home/proja/proj1rpt
```

The **dsmrecall** can also be invoked directly, by using a shell application or a script to build a list of files to be recalled. This file list can then be passed directly to HSM. For example, the following command recalls all files owned by user ibm:

```
find /hsmmanagedfilesystem -user ibm -print > /tmp/filelist  
dsmrecall -filelist=/tmp/filelist
```

Related concepts:

“Recalling migrated files overview” on page 3

Related reference:

“**dsmrecall**” on page 157

Normal recall mode

Normal recall mode is the default for all files. Files are recalled completely from IBM Spectrum Protect server storage and can be accessed after the recall process is complete.

Partial file recall mode

In a GPFS environment, the space management client provides a partial file recall mode to recall a portion of a migrated file. This avoids having to recall an entire, potentially large file, when only a small portion of the file is required by an application. This recall mode is only for read access on migrated files that were transferred without compression. Files that were migrated by using compression are always recalled completely. Partial file recall mode cannot be used on executable files.

If a file is larger than the value of the **minpartialrecallsize** option and compression is turned on, the partial file recall mode prevails, and the file is migrated without compression.

When a vendor application makes a read request for a file that is qualified for partial file recall, and the file is migrated, the space management client calculates which portion of the file to recall based on the offsets contained in the read request. This results in time and disk space savings, because only a portion of the file is recalled.

When you use partial file recall and need to work with a large portion of a large file that has been migrated, it is possible to recall a specified portion of the file in one operation. This minimizes the number of partial file recall requests and access

to the server storage based on the **minpartialrecallsize** option value. Use the **dsmrecall** command with the **offset** and **size** options to specify the data range within the file.

For example, to recall 200 MB of the file `/usr/cam/video2`, starting at 400 MB from the beginning of the file, issue the following command:

```
dsmrecall -offset=400m -size=200m /usr/cam/video2
```

You can use the following methods to specify which files HSM should recall using partial file recall:

- Set the **minpartialrecallsize** option of the **dsmmigfs** command to the minimum size (in megabytes) that a file must have to qualify for partial file recall.
- Set the **recallmode** option of the **dsmattr** command to `partialrecall`. This specifies that, regardless of its size, a file is recalled using partial file recall.

Use this method to change the recall mode of migrated files that you normally read but you do not modify. When you set the recall mode to partial file recall, this mode remains associated with your files until you:

- Change the recall mode
- Modify the file
- Recall the file selectively
- Restore the file

Related concepts:

“Recalling migrated files overview” on page 3

Related reference:

“**dsmmigfs add** and **update**” on page 128

“**dsmattr**” on page 118

“**dsmrecall**” on page 157

Streaming recall mode

Streaming recall mode allows for an asynchronous recall of migrated files. The recalled portion of the file can be accessed while the file is recalled.

Note:

- Streaming recall mode is valid only for read-only operations on the file.
- Partial file recall mode takes precedence over streaming recall mode. If a file is smaller than the value of **minpartialrecallsize** or if the **minpartialrecallsize** option is set to 0, normal or streaming recall mode takes precedence.

You can use the following methods to specify which files are recalled with streaming recall mode:

- Set the **recallmode** option of the **dsmattr** command to `s` (streaming). This option specifies that you want to enable an asynchronous recall of migrated files.

Use this method to change the recall mode of migrated files that you typically read but that you do not modify. Streaming recall mode remains associated with the files until you do any of the following actions:

- Change the recall mode
- Modify the file
- Recall the file selectively
- Restore the file

- Set the **minstreamfilesize** option of the **dsmmigfs add** or **dsmmigfs update** command to specify a number to enable or disable an asynchronous recall of migrated files.

A larger value of the **streamseq** option of the **dsmmigfs add** or **dsmmigfs update** command increases the performance of streaming recall.

The **readeventtimeout** option of the **dsmmigfs add** or **dsmmigfs update** command sets the time before a streaming recall process times out.

The **readstartsrecall** option of the **dsmmigfs add** or **dsmmigfs update** command specifies whether a recall operation starts immediately when an application reads the stub file.

Related concepts:

“Recalling migrated files overview” on page 3

Related reference:

“**dsmattr**” on page 118

“**dsmmigfs add** and **update**” on page 128

How HSM determines which recall mode to use

The space management client determines the recall mode depending on the configured recall mode and whether a file is modified.

Table 13 displays the recall mode that is used. The recall mode depends on the following factors:

- The recall mode that you configured for a migrated file
- Whether the file is modified

Table 13. space management client determining recall modes

The configured recall mode:	The file is modified:	The recall mode and file state:
Normal	No	Normal recall mode. The file becomes premigrated.
Normal	Yes	Normal recall mode. The file becomes resident.
Streaming	No	Streaming recall mode. The file becomes premigrated.
Streaming	Yes	Normal recall mode. The file becomes resident.
Partial	No	Partial recall mode. The file stays migrated.
Partial	Yes	Normal recall mode. The file becomes resident

Setting the recall mode by using the **dsmattr** command

To set or change the recall mode for one or more migrated files, use the **dsmattr** command. Select normal, partial recall, or streaming recall mode.

Procedure

Enter the **dsmattr** command. Identify the files and the recall mode by using the command parameters.

Related reference:

“**dsmattr**” on page 118

Optimized tape recall processing

When files reside on tape, the space management client can order a list of files to recall the files efficiently. Optimized tape recall processing is a feature of the selective recall function.

When files are on tape, it can take a long time to recall a list of files. The requested files can be on different tapes, and the files in the list might not be grouped in the same way as the files are stored on tapes. Tapes can be frequently mounted and unmounted. You can optimize tape access and minimize mounting and unmounting tapes.

To optimize tape recall processing, the **dsmreca11** command creates ordered lists of files. The command creates one ordered list per tape for files that are stored on tape, and one list for all files that are stored on disk media. Each tape list is ordered to recall the files from tape in the most efficient order.

Optimized tape recall processing provides the following benefits:

- Optimized processing avoids frequent tape mount and unmount operations.
- Optimized processing avoids excessive tape seek operations.
- If the files are located contiguously on tape, optimized processing enables streaming mode.
- Optimized processing orders and recalls files in separate steps so that you can restrict the recall operations on certain tapes.
- Optimized processing can recall files from several tape drives in parallel to increase recall throughput.

The user ID that you use to optimize tape recall processing requires root permission.

Optimized tape recall is a feature of selective recall. Optimized tape recall is not a feature of transparent recall.

Contention for tape resources

Transparent recall processing does not decrease the efficiency of the selective, optimized tape recalls. But the two kinds of recall processing must wait for each other when requesting the same resources. For example, if a transparent recall process requests a tape drive that is being used by an optimized tape recall process, the transparent recall process waits. When the optimized tape recall process is finished with the tape drive, the transparent recall process can use the tape drive.

You can limit the tape drives available for tape optimized processing with the **hsmmaxrecalltapedrives** option. Limiting the tape drives available for tape optimized processing frees tape drives for other operations like selective recall or migration.

Related concepts:

“Recalling migrated files overview” on page 3

Related reference:

“**dsmreca11**” on page 157

“**hsmmaxrecalltapedrives**” on page 101

List files for optimized tape recalls

You can preview a list of files for optimized tape recalls without starting the recall process.

The **preview** option with the **dsmrecall** command generates the list files without starting the recall process. Each run with the **preview** option generates a set of files in a different directory. You can run simultaneous **dsmrecall** processes on the same node or on different nodes with different input list files. You can run simultaneous **dsmrecall** processes for different file systems or for the same file system.

The **dsmrecall** command generates the list files every time that it runs, even if you do not specify the **preview** option. The generated file lists are not deleted automatically. You must delete the list files when you no longer require them.

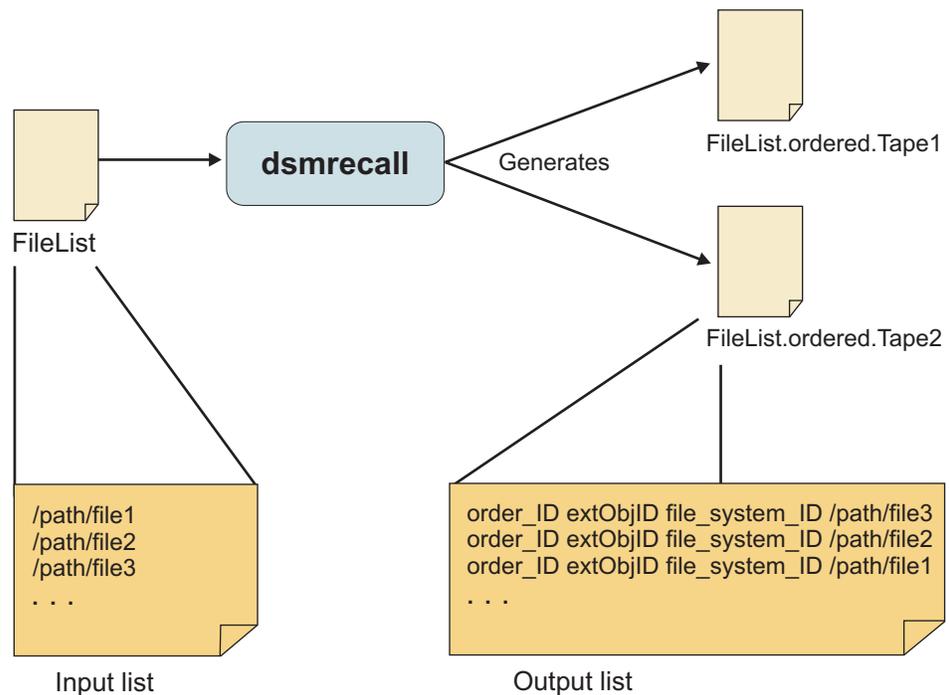


Figure 2. Ordered list file generation

Input file lists

In addition to the rules specified for list files, the files listed in the input list file must follow this rule:

- Each entry in the list file must be from the same file system, and you must specify the file system when you start the **dsmrecall** command.

Any files from other file systems are skipped and written to a list of unprocessed files. This list of unprocessed files is stored in `.SpaceMan/tapeOptimizedRecall/FileList.unprocessed.file_system_name`. The list file of unprocessed files can be used as a new input list file for the next **dsmrecall** command.

Generated list files

There are two types of generated list files. All generated files are stored in `.SpaceMan/tapeOptimizedRecall` in the space-managed file system. Each time `dsmrecall` creates ordered list files, they are created in a new directory. You can run several processes simultaneously without overwriting the generated files. The directory includes the node ID (*node_ID*) and process ID (*PID*): `.SpaceMan/tapeOptimizedRecall/node_ID/PID`.

Tape and disk list files

One ordered list file is generated for each tape. The tape list file has the following properties:

- The tape list file has a name with this format: `filelist.ordered.tape.tapeID`.
- Each entry in the tape file has this format:
`order_ID extObjID file_system_ID /filesystem/path/filename`
- Files in these lists are ordered by their position on the tape.

There is one file list for all files stored on a disk pool on the server. The disk list file has the following properties:

- The disk list file has a name with this format: `filelist.nonTape.00000000`.
- Each entry in the disk file has a format with this format:
`order_ID extObjID file_system_ID /filesystem/path/filename`
- Files in the disk list are not ordered since they are on disk.
- You cannot specify a tape list file or a disk list file on the `filelist` option of the `dsmrecall` command.

Collection files

There is one collection file that includes the names of the tape list files and the disk list file. The collection file has these properties:

- The collection list file has a name with this format:
`filelist.ordered.collection`

Entries in this list are the names of the tape and disk list files.

- Entries in the collection file have this format:
`filecount byteCount filelist.nonTape.00000000`
`filecount byteCount filelist.ordered.tape tapeID1`
`filecount byteCount filelist.ordered.tape tapeID2`

filecount is the number of files in the list, and *byteCount* is the space that is occupied in the system when all files from the list are recalled.

- By default the collection file is ordered by the tape ID in ascending order. The disk list file is in the first position. The recall is done in order of the entries, thus the first entry has the highest recall priority, and last entry has the lowest recall priority. You can manually rearrange the order to suit your environment.
- You can specify a collection file on the `filelist` option of the `dsmrecall` command.

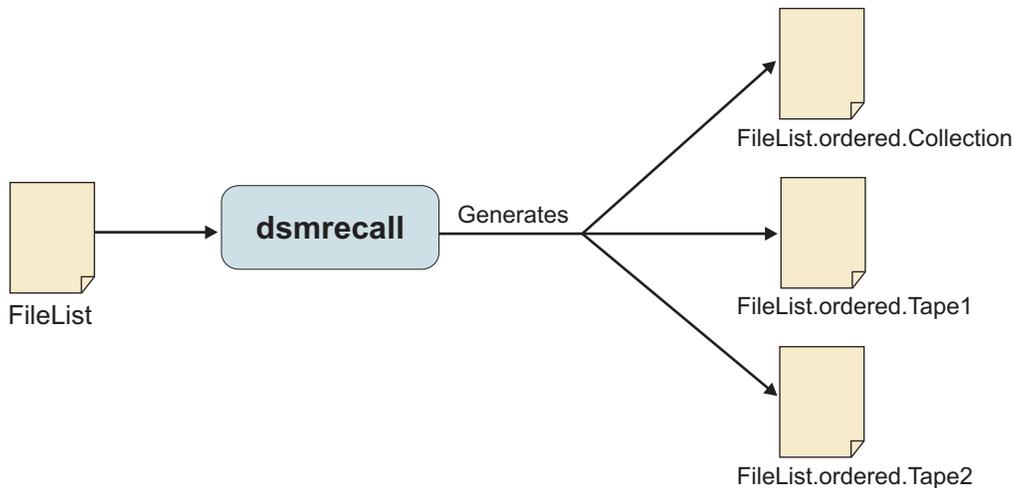


Figure 3. Generation of list files

Related concepts:

Chapter 11, “Options files reference,” on page 89

Related reference:

“**dsmrecall**” on page 157

Recalling files with optimization for tape

Tape processing is automatically optimized when you invoke **dsmrecall** with the **filelist** option and specify a file system.

About this task

If you do not specify a file system, the recall process does not optimize tape processing.

Procedure

Start a selective recall with the **filelist=list_file** option and specify a file system. For example:

```
dsmrecall -filelist=myFileList myFileSystem
```

If the list file is a collection file that was generated by the **dsmrecall** command with the **preview** option, the recall begins immediately. The files are recalled in the order specified by the collection file.

If the list file is not a collection file that was generated by the **dsmrecall** command, and the list file is appropriately formatted, the file entries are sorted for optimized tape processing and are then recalled.

The recall begins immediately in the following example, because the file identified by the **filelist** parameter is a collection file.

```
dsmrecall -filelist=/HsmManagedFS/.SpaceMan/tapeOptimizedRecall/node_ID/PID/
FileList.ordered.collection myFileSystem
```

Related concepts:

“List files for optimized tape recalls” on page 76

Related reference:

“**dsmrecall**” on page 157

Chapter 8. File system reconciliation

To keep local file systems synchronized with the IBM Spectrum Protect server for space management services, the space management client automatically reconciles your file systems at preset intervals. You can also start reconciliation manually, and you can configure immediate reconciliation.

Note: To reconcile file systems, you must have root user authority.

Related concepts:

Chapter 1, “Overview of the space management client,” on page 1

“Reconciling file systems overview” on page 4

Automatic reconciliation

The space management client automatically reconciles each file system for which space management is active. For example, when you modify or delete a migrated or premigrated file from your local file system, an obsolete copy of the file remains in storage. During automatic reconciliation, any obsolete copies of your migrated or premigrated files are marked for expiration.

To specify how often reconciliation runs, modify the setting on the **reconcileinterval** option in your `dsm.sys` options file. The default is every 24 hours.

You can specify how many days a migrated or premigrated file remains in IBM Spectrum Protect server storage after you recall and modify or erase it from your local file system. Modify the setting on the **migfileexpiration** option in your `dsm.sys` options file. The default is 7 days. When the copies expire, they are removed from the server.

Table 14 describes the tasks that automatic reconciliation completes for files and file systems.

Table 14. Automatic reconciliation tasks

File type	Reconciliation Tasks
Migrated files	<ul style="list-style-type: none">• Verifies that a stub file exists on your local file system for each migrated file in storage• Marks a migrated file for expiration in the following cases:<ul style="list-style-type: none">– You deleted the stub file from your local file system– You recalled the file and modified it• Removes a migrated file from IBM Spectrum Protect server storage if it expired• Updates the status file
Pre migrated files	<ul style="list-style-type: none">• Verifies that premigrated files are still valid• Marks a premigrated file for expiration in the following cases:<ul style="list-style-type: none">– You deleted the file from your local file system– You accessed the file and modified it• Removes a copy of a premigrated file from IBM Spectrum Protect server storage if it expired.• Updates the status file

Table 14. Automatic reconciliation tasks (continued)

File type	Reconciliation Tasks
Stub files	Records the name of any file for which a stub file exists on your local file system, but a migrated file does not exist in storage. The names are recorded in the <code>orphan.stubs</code> file.
Status file	<p>Updates the following information in the status file:</p> <ul style="list-style-type: none">• Number of premigrated files• Number of premigrated blocks• Number of migrated files• Number of migrated blocks <p>Note: Migration and recall processes update status information dynamically. If any other process changes the state of a file, the status file does not reflect the change until reconciliation is run.</p> <p>Related concepts: Chapter 11, “Options files reference,” on page 89 “Orphaned stub files” on page 81</p>

Manual reconciliation of file systems

If you recall a migrated file, modify it, and selectively migrate it to IBM Spectrum Protect server storage, two copies of that file reside in storage. The unmodified copy of the file now is obsolete. Manual reconciliation synchronizes your file system and saves space by removing obsolete copies from IBM Spectrum Protect server storage.

If you set the `migfileexpiration` option to 0 in your `dsm.sys` options file, you can run reconciliation immediately to delete the obsolete copies from IBM Spectrum Protect server storage and create available space for your migrated files.

Important: A value of 0 indicates that an obsolete copy of a migrated or premigrated file is directly deleted from the server during the next reconciliation run. If you delete the file from the local file system and the reconciliation process runs with the `migfileexpiration` option set to 0, the file cannot be re-created by the `dsmmigundelete` command.

The `migfileexpiration` option is also used by the automatic reconciliation process.

You can also use the `dsmreconcile` command to search for orphan stub files and complete metadata updates on the server. The following example for the `/home` file system shows how to search for orphan stub files:

```
dsmreconcile -o /home
```

Note:

- Reconciliation requires that the scout daemon is running.
- If you reconcile several file systems, increase the value on the `reconcileinterval` option in the `dsm.sys` file to reduce the impact that the `dsmreconcile` command might have on system performance.

After you run reconciliation, check the `orphan.stubs` file in the `.SpaceMan` directory for each file system that you reconciled. The file lists orphaned stub files.

Related concepts:

“The scout daemon” on page 84

“Orphaned stub files”

Related reference:

“**dsmreconcile**” on page 161

Orphaned stub files

An orphaned stub file is a stub file for which a corresponding migrated file in IBM Spectrum Protect server storage is not located. If orphaned stub files exist in your file systems, the space management client records information about these files in the `orphan.stubs` file during reconciliation.

If you set the **errorprog** option in your `dsm.sys` file, a message is sent to the program that you specified with this option during automatic reconciliation.

To check for orphaned files, specify YES on the **checkfororphans** option in the `dsm.sys` file. When orphaned files are located, their names are recorded in the `.SpaceMan/orphan.stubs` file. If you specify YES, the **dsmreconcile** process queries the scout daemon for all migrated and premigrated files and checks that corresponding objects exist on the server.

Stub files might become orphaned include the following situations:

- You modified your `dsm.sys` file so your client node now contacts a different server for space management services other than the one to which files were migrated.
 - To resolve this problem, modify your `dsm.sys` file so your client node contacts the server to which the files migrated.
- Your IBM Spectrum Protect administrator uses the **delete filespace** administrator command to delete any migrated files from a specific file system.
 - To resolve this problem, if files are no longer needed, an administrator can delete some or all of the files from IBM Spectrum Protect server storage. The stub files are no longer valid and you can erase them.
- A media failure occurs that corrupts or loses your migrated files. Storage pool backup and recovery provides protection against media failures. However, if you cannot restore a migrated file from a migration storage pool, you can restore a backup version of the file if you used the backup-archive client.
 - When you set the **restoremigstate** option to NO in your `dsm.opt` file, and you then restore a backup version of a migrated file, the file becomes a normal, resident file.

Configuring immediate reconciliation

The space management client can expire migration copies on the IBM Spectrum Protect server when migrated or premigrated files are deleted from the space-managed file system.

About this task

You can configure immediate reconciliation. When a migrated or premigrated file is deleted on the space-managed file system, the migration copy on the IBM Spectrum Protect server is expired immediately. Expirations are synchronized with deletions when a file system is space managed by a single IBM Spectrum Protect server or multiple IBM Spectrum Protect servers.

If you set `MIGFILEEXPIRATION 0`, the migration copy on the IBM Spectrum Protect server is deleted immediately. You cannot recover the file on the space-managed file system with the `dsmmigundelete` command.

A separate reconciliation is required to expire orphans on the IBM Spectrum Protect server if the process is interrupted by a system failure. A separate reconciliation is required to expire orphans on the IBM Spectrum Protect server if the process is interrupted by a failover of space management. A separate reconciliation is required to update metadata information of migrated and premigrated files on the IBM Spectrum Protect server. The metadata includes path and file name, time stamps, owner, group and access rights.

Procedure

1. Set `HSMEVENTDESTROY yes` in the `dsm.opt` options file.
2. Stop and start again all HSM daemons.

Reconciling by using a GPFS policy

You can use the General Parallel File System (GPFS) policy engine to reconcile a file system with the corresponding IBM Spectrum Protect server. Use the GPFS policy engine to reconcile a file system that is managed by one or several IBM Spectrum Protect servers.

Before you begin

You must configure HSM-GPFS integration before you use the `dsmreconcileGPFS.pl` command. The integration must include a GPFS policy that identifies all migrated and premigrated files.

About this task

The HSM GPFS client can use the GPFS policy engine to reconcile file systems with IBM Spectrum Protect servers. In a multiple-server environment, you must specify the IBM Spectrum Protect server to reconcile.

Procedure

To reconcile a file system, issue the following command: `dsmreconcileGPFS.pl file_system_name`.

In a multiple-server environment you must specify a server to reconcile with the **server** option: `dsmreconcileGPFS.pl -server=server_name file_system_name`. To reconcile the file system with all servers, start the `dsmreconcileGPFS.pl` script with each IBM Spectrum Protect server.

The `dsmreconcileGPFS.pl` script uses the GPFS policy engine to determine which files are migrated. The script also starts the `ds mreconcile` command to complete a two-way reconciliation by using the **orphancheck** and **fileinfo** options.

Related reference:

“`dsmreconcileGPFS.pl`” on page 170

“`ds mreconcile`” on page 161

Chapter 9. Space management daemons

The space management client uses the space monitor, recall, scout, and watch daemons to manage file systems automatically.

These daemons are included with the space management client. The space monitor daemon (**dsmonitor**), the scout daemon (**dsmscoutd**), and the recall daemon (**dsmrecalld**) manage space on GPFS file systems. The watch daemon (**dsmwatchd**) runs only on GPFS file systems.

The daemons start when you add space management to your file systems and when you modify space-management options.

On GPFS file systems, you can use the GPFS policy engine to monitor space thresholds, search for migration candidates, and identify migrated files for reconciliation. If you use the GPFS policy engine, you can disable the space monitor and scout daemons.

For information about configuring GPFS integration with the space management client, see Technote 7018848.

The space monitor daemon

The space monitor daemon monitors space usage on all file systems and the storage pools in those file systems to which you add space management. It starts threshold migration whenever necessary.

To check space usage more frequently or less frequently, change the value on the **checkthresholds** option in the `dsm.sys` file. To reconcile your file systems more frequently or less frequently, change the value on the **reconcileinterval** option in the `dsm.sys` file. When you change the option values the space monitor daemon uses, the new values are not effective until you stop and restart the space monitor daemon.

The space monitor daemon starts automatically when you mount the file system and add space management to it. If the space monitor daemon stops running, issue the **dsmonitor** command to start it.

When GPFS policy drives automatic migration, disable the space monitor daemon by setting the option **hsmdisableautomigdaemons=YES**.

For information about configuring GPFS integration with the space management client, see Technote 7018848.

Related reference:

“**reconcileinterval**” on page 108

“**checkthresholds**” on page 91

“**dsmonitor**” on page 155

“**hsmdisableautomigdaemons**” on page 94

The recall daemon

The recall daemon recalls migrated files from storage to the local file system. The "master" recall daemon starts two more instances by default that do not recall files: "distributor" and "receiver". These instances are used for communication purposes only.

If a recall daemon is not running, enter the **dsmrecalld** command to start one.

A child recall daemon can recall only one file at a time; however, you can run more than one recall daemon at the same time. To set the minimum and maximum number of child recall daemons that you want to run at one time, use the **minrecalldaemons** and **maxrecalldaemons** options in the `dsm.sys` file. The minimum number of child recall daemons that you can run at the same time is 1. The default is 3. The maximum number of child daemons that you can run at the same time is 99. The default is 20. On a General Parallel File System (GPFS) the maximum is constrained by the GPFS option **dmapiWorkerThreads**.

If all child recall daemons are busy, another file cannot be recalled until a child recall daemon is available. If an application uses all available child recall daemons because all files are migrated, the application is suspended until a child recall daemon is available. You can increase the value of the **maxrecalldaemons** option.

When you change the option values that the recall daemons use, the new values are not effective until you stop and restart the recall daemons.

Related reference:

"**maxrecalldaemons**" on page 105

"**minrecalldaemons**" on page 108

"**dsmwatchd**" on page 166

"**dsmrecalld**" on page 160

The scout daemon

The scout daemon automatically searches for candidates on each file system or storage pool in a file system for which space management is active.

The scout daemon scans file systems and stores the information for each file in a complete file index (CFI). The daemon works with the CFI to search for migration candidates. The CFI is updated automatically during all migration, recall, and restore operations.

To specify how often the scout daemon scans a file system, modify the setting on the **candidatesinterval** option in the `dsm.sys` file.

When GPFS policy drives automatic migration, disable the scout daemon by setting the option **hsmdisableautomigdaemons=YES**.

For information about configuring GPFS integration with the space management client, see Technote 7018848.

Related reference:

"**dsmscoutd**" on page 164

"**candidatesinterval**" on page 90

The watch daemon

The watch daemon checks the status of other HSM daemons. If any of the other daemons end or become corrupted, the watch daemon automatically recovers the failed daemon.

The watch daemon checks the status of the recall, the monitor, and the scout daemons.

On Red Hat Enterprise Linux Version 6 (RHEL6), the **dsmwatchd** daemon is started from the **initctl** service. It can be started manually with the following command: **initctl start HSM**. Stop the watch daemon with the following command: **initctl stop HSM**.

On all other systems besides RHEL6, you cannot start the watch daemon manually. It is started by the **init** service.

Related reference:

“**dsmwatchd**” on page 166

Stopping the space management daemons

Use the following steps to properly stop space management daemons.

About this task

Do not use the **kill -9** command to stop any space management daemon. All daemons have their own clean-up procedure, which is interrupted if you use the **kill -9** command. Using the **kill -9** command can yield unpredictable and unintentional results.

Follow these steps to properly stop the space monitor, master recall, subordinate recall, or scout daemon.

Note: Stopping a master recall daemon stops all subordinate daemons.

Procedure

1. Issue the **dsmq** command to obtain the recall ID and the recall daemon process ID for each recall process that is in the queue.
2. Issue the **dsmrm** command to remove each recall process from the queue.
3. Issue **ps -ef | grep dsm** to verify that both the space monitor daemon and the master recall daemon are running.
4. Issue the **kill -15** command with the process identifier number to stop the daemons.
5. On GPFS file systems, issue the command **dmkilld** to stop the recall daemons.
6. Verify that the daemons are no longer running. For AIX and Linux GPFS file systems, issue the **dsmmigfs stop** command to stop all space management daemons.
7. Stop the **dsmwatchd** daemon.

On Red Hat Enterprise Linux Version 6 (RHEL6) systems:

Stop the daemon with the following command: **initctl stop HSM**.

On all other systems besides RHEL6:

- a. Comment out the entry for the **dsmwatchd** daemon in the file `/etc/inittab`.
- b. Close the file and submit the change by issuing the following command: `telinit Q`.
- c. Stop the **dsmwatchd** daemon with the command `kill -15`.

Related reference:

"**dsmq**" on page 156

"**dsmrm**" on page 163

"**dmkilld**" on page 118

"**dsmmigfs stop, start, and restart**" on page 144

Chapter 10. Scheduled HSM service

IBM Spectrum Protect provides a service to schedule backup, archive, and space management tasks.

Scheduling requires a cooperative effort between the IBM Spectrum Protect server and your client node.

Your IBM Spectrum Protect administrator defines a schedule on the server and associates your client node with that schedule to perform backup, archive, or space management tasks automatically. The administrator sets server parameters for the following schedule tasks:

- Balance scheduled services for all client nodes
- Specify that your client node can query the server for scheduled work at specific time intervals or wait for the server to contact your client node
- Control how often your client node contacts the server for scheduled work

Before scheduled services can be started, you must set scheduling options in your `dsm.sys` file and start a client scheduler on your workstation. The backup-archive command-line client must be installed to start the client scheduler.

The IBM Spectrum Protect backup-archive client **schedule** command starts the client scheduler.

The backup-archive client **query schedule** command displays information about scheduled services for your client node.

You are not required to use IBM Spectrum Protect to schedule space management services. You can set up a cron job to run space management services at specific times. If you set up a cron job, set the **reconcileinterval** option to 0 in your `dsm.sys` file so the space management client does not automatically reconcile file systems at specific intervals.

Related information:

-  [Query schedule client command](#)
-  [Schedule command](#)
-  [Schedule operations for backup-archive clients](#)

Chapter 11. Options files reference

The space management client provides system and space management options that you set either in your `dsm.sys` file or in your `dsm.opt` file. The values that you set for space management options determine which server your client node contacts for space management services and affect automatic migration, reconciliation, and recall.

Related tasks:

“Editing the options file `dsm.opt`” on page 22

“Editing the options file `dsm.sys`” on page 20

afmskipuncachedfiles

The **afmskipuncachedfiles** option specifies whether uncached and dirty files in General Parallel File System (GPFS) Active File Management file sets are processed for backup, archive, and migration operations.

GPFS Active File Management and *uncached* and *dirty* file states are explained in General Parallel File System product information.

Running HSM on GPFS file systems that use Active File Management file sets is explained in *Configuring IBM Spectrum Scale Active File Management*. If you back up, archive, or migrate files from a file system that contains Active File Management file sets, set `afmskipuncachedfiles=yes`.

Restriction: If Active File Management is running in Local Update (LU) mode, the **afmskipuncachedfiles** option in the cache file set must be set to **No**.

Options File

Place this option at the beginning of the `dsm.sys` file before any server stanzas.

Syntax



Parameters

NO The Active File Management file state is ignored during backup, archive, and migration operations. Migration operations on uncached or dirty files fail and yield error message ANS9525E. Back up and archive operations on uncached files require Active File Management fetch operations. The fetch operations can cause significant network traffic between the Active File Management home and cache.

YES

Uncached or dirty files in Active File Management file sets are skipped during backup, archive, and migration processing.

candidatesinterval

The **candidatesinterval** option specifies how often the **dsmscoutd** daemon searches the file systems for migration candidates.

The **dsmscoutd** daemon scans the file systems and stores the information in a complete file index (CFI), which is used to search for migration candidates.

Tip: When automatic migration is driven by GPFS policy, the **dsmscoutd** daemon is disabled, and the **candidatesinterval** option has no effect.

Options File

Place this option at the beginning of the `dsm.sys` file before any server stanzas.

Syntax

►—CANDIDATESINTERVAL—*scan_interval*—◄

Parameters

scan_interval

Specifies the maximum interval of time that elapses between each successive time the **dsmscoutd** daemon automatically scans for files in the file systems. The range of values is 0 to 9999. The default is 1.

Specify 0 to continuously scan the file systems. When scanning reaches the end of a file system, the **dsmscoutd** daemon immediately begins scanning again from the beginning of the file system.

Specify 1 to scan the file systems at intervals that depend on the percentage of the file system content that has changed. The **dsmscoutd** daemon increases the frequency of scanning as the percentage of file system changes increases. The **dsmscoutd** daemon reduces the frequency of scanning as the percentage of file system changes decreases. This is the default.

Specify a value from 2 to 9999 to define the number of hours between file system scans. After scanning reaches the end of the file system, the **dsmscoutd** daemon waits the specified number of hours before starting the next scan.

Related reference:

“**dsmscoutd**” on page 164

checkfororphans

The **checkfororphans** option specifies whether or not the **dsmreconcile** command checks for migrated files that are no longer present on the server but whose stub files still remain on the client (orphans). The option parameter that you set determines whether or not the **dsmreconcile** command queries the scout daemon for all migrated and premigrated files.

Options File

Place this option at the beginning of the `dsm.sys` file before any server stanzas.

Syntax

►►—CHECKFororphans—

NO
YES

—————►►

Parameters

NO The **dsmreconcile** command checks for orphans only when you remove the space management client from your file system, or when you specify the **-o** parameter with the **dsmreconcile** command. Queries to the scout daemon for migrated and premigrated files are not done. This is the default.

YES

The **dsmreconcile** command checks for orphans. When orphans stubs are located, their names are recorded in `.SpaceMan/orphan.stubs`. The **dsmreconcile** command queries the scout daemon for all migrated and premigrated files.

The scout daemon must be running if the **checkfororphans** option is set to YES.

Related reference:

“**dsmreconcile**” on page 161

“**dsmscoutd**” on page 164

checkthresholds

The **checkthresholds** option specifies how often the space monitor daemon checks space usage on your file systems. The space monitor daemon checks each file system to which you added space management.

Tip: When automatic migration is driven by GPFS policy, the **dsmmonitord** daemon is disabled, and the **checkthresholds** option has no effect.

Options File

Place this option at the beginning of the `dsm.sys` file before any server stanzas.

Syntax

►►—CHECKThresholds—*interval*—————►►

Parameters

interval

Specifies the number of minutes that must elapse before the space monitor daemon checks space usage on your file systems. The range of values is 1 through 9999. The default is 5.

Related reference:

“**dsmmonitord**” on page 155

compression

The **compression** option compresses files before you send them to the server. Compressing your files reduces data storage for backup versions and archive copies of your files. It can, however, affect IBM Spectrum Protect throughput.

Tip: This option controls compression only if your administrator specifies that your client node determines the selection. The server also can define this option.

A fast processor on a slow network connection benefits from compression, but a slow processor on a fast network connection does not. Use this option with the backup-archive client option, **compressalways**.

If you specify **compressalways YES**, compression continues even if the file size increases. To stop compression if the file size grows, and resend the uncompressed file, specify **compressalways NO**.

If you specify **compression YES**, you can control compression processing in the following ways:

- Use the **exclude.compression** option in your include-exclude options file to exclude specific files or groups of files from compression processing.
- Use the **include.compression** option in your include-exclude options file to include files within a broad group of excluded files for compression processing.

Options File

Place this option in the client system options file `dsm.sys` within a server stanza.

Syntax



Parameters

NO Files are not compressed before they are sent to the server. This is the default.

YES

Files are compressed before they are sent to the server.

defaultserver

The **defaultserver** option specifies the default server to which you back up and archive your files from your local file systems. If you do not specify a migration server with the **migrateserver** option, this option can also specify the server to which files are migrated from your local file systems.

Options File

Place this option at the beginning of the `dsm.sys` file before any server stanzas.

Syntax

►►—DEFAULTServer—*servername*—◄◄

Parameters

servername

Use the **defaultserver** option to specify the name of the IBM Spectrum Protect server to contact for backup-archive services if more than one server is defined in the `dsm.sys` file. By default, the space management client contacts the server that is defined by the first stanza in the `dsm.sys` file. This option is only used if the *servername* option is not specified.

If you do not specify a migration server with the **migrateserver** option, this option specifies the server to which you want to migrate files.

You can override this option with the following command: `dsmmigfs upd /FS -Server=servername`. Replace *servername* with the name of your server.

The value of **defaultserver** in the `dsm.sys` file overrides **defaultserver** in the `dsm.opt` file.

Related reference:

“**migrateserver**” on page 106

errorlogname

The **errorlogname** option specifies the path of the file that logs error messages for the space management client and the backup-archive client.

This option specifies the fully qualified path and file name of the error log file. If this option is not specified, the default log file is used. The default error log file is the `dsmerror.log` file in the current working directory.

Authorized user: The value for this option overrides the **DSM_LOG** environment variable.

The log file path cannot be a symbolic link.

Options File

Place this option in the client system options file `dsm.sys` within a server stanza.

Syntax

►►—ERRORLOGName—*filespec*—◄◄

Parameters

filespec

The fully qualified path in which to store error log information. If any part of the path you specify does not exist, the space management client creates it.

The log file path cannot be a symbolic link.

errorprog

The **errorprog** option specifies a program to which you want to send a message if a severe error occurs during space management processing.

The **errorprog** option specifies a program to which you want to send a message if a severe error occurs during space management processing.

The program is started only in case of very severe errors such as DMI or file system errors. Also, only processes that do not have an output channel, such as daemons, use this option. Use the full path to specify the program. The error message is sent to the standard input of the program.

Options File

Place this option at the beginning of the `dsm.sys` file before any server stanzas.

Syntax

▶—`ERRORProg—program-name`—▶

Parameters

program-name

Specifies the path and file name of the program to which you want to send a message if a severe error occurs during space management processing.

Some examples:

```
errorprog /usr/bin/cat >/tmp/tsm_severe_errors.txt
errorprog /usr/bin/perl /root/dsmsevererror.pl
```

hsmdisableautomigdaemons

Use the **hsmdisableautomigdaemons** option to control the start of the **dsmscoutd** and **dsmmonitord** daemons.

Supported Clients

When a GPFS policy monitors migration thresholds, the **dsmscoutd** and **dsmmonitord** daemons can be disabled. Stop and restart the **dsmwatchd** daemon to activate a new value of the **hsmdisableautomigdaemons** option. The **dsmwatchd** daemon is restarted by the **init** process or the **initctrl** process.

This option is valid only on GPFS file systems.

Options File

Place this option in the client options file `dsm.opt`.

Syntax

▶—`HSMDISABleautomigdaemons`—NO
YES—▶

Parameters

NO Specifies that the space management client automatic migration daemons **dsmscoutd** and **dsmonitord** are started. This is the default.

YES

Specifies that the **dsmscoutd** and **dsmonitord** daemons are not started.

Tip: The **dsmscoutd** and **dsmonitord** daemons can also be started by the GPFS daemon (**mmfsd**).

Related tasks:

“Stopping the space management daemons” on page 85

hsmdistributedrecall

The **hsmdistributedrecall** option specifies whether a node that is running as owner node can delegate recall requests to other nodes in a GPFS cluster.

To stop distributed recalls for an entire cluster, set this option to **NO** for all HSM nodes in the cluster and restart the **dsmd** daemons. Then, only the HSM owner node can recall migrated files to the space-managed file systems it owns.

Options File

Place this option at the beginning of the `dsm.sys` file before any server stanzas.

Syntax



Parameters

YES

The node can delegate recall requests to other nodes in a GPFS cluster. This is the default.

NO The node cannot delegate recall requests to other nodes in a GPFS cluster.

hsmenableimmediatemigrate

The **hsmenableimmediatemigrate** option determines if files that are less than 2 minutes old can be migrated during selective migration.

When a file is created, file status information can take some time to become synchronized among GPFS nodes. If you create a file, then migrate the file before file information is synchronized, the file size information returned from a **dsmd** or **dsmdf** command can be inaccurate. Because of the typical time needed for GPFS synchronization, by default the space management client does not migrate a file that is less than 2 minutes old.

To migrate files that are less than 2 minutes old, set **hsmenableimmediatemigrate=YES**. When you set **hsmenableimmediatemigrate=YES**, the space management client issues system synchronization commands before migrating each file that is less than 2 minutes old.

This option applies only during selective migration.

Options File

Place this option in the client options file `dsm.opt`.

Syntax



Parameters

NO Files that are less than 2 minutes old cannot be migrated. This is the default.

YES

Files that are less than 2 minutes old can be migrated during selective migration.

Related reference:

“`dsmmigrate`” on page 149

hsmeventdestroy

The **hsmeventdestroy** option specifies whether immediate reconciliation occurs.

When you set `HSMEVENTDESTROY` yes, the `dsmreconciled` daemon starts on the affected nodes. The `dsmreconciled` daemon synchronizes the expirations on the IBM Spectrum Protect server with deletions on the space-managed file system.

Options File

Place this option in the client options file `dsm.opt`.

Syntax



Parameters

NO Immediate reconciliation does not occur. A migrated or premigrated file that is deleted on the space-managed file system is expired on the IBM Spectrum Protect server only during a reconciliation of the file system. This is the default.

YES

Immediate reconciliation occurs. When a migrated or premigrated file is deleted on the space-managed file system, the migration copy on the IBM Spectrum Protect server is expired immediately.

hsmextobjidattr

The **hsmextobjidattr** option specifies whether the **dsmreconcile**, **dsmmigrate**, **dsmrecall**, or **dsmrecalld** command assigns an object ID to the file as a Data Management Application Programming Interface (DMAPI) extended attribute (extObjId).

The extObjId DMAPI extended attribute is required for two-way orphan check reconciliation.

The two-way orphan check processing identifies orphans on the file system and on the IBM Spectrum Protect server. The two-way check is done in a single pass and the orphan identification process uses parallel processing.

Place this option in the client options file `dsm.opt`.

Syntax



Parameters

NO An object ID is not assigned to the file that is processed by the **dsmmigrate**, **dsmrecall**, **dsmrecalld**, or the **dsmreconcile** command with the **preptwo** option. When a reconciliation is run with the **preptwo** option, files are not prepared for the two-way orphan check. This is the default.

YES

An object ID is stored in the file that is processed by the **dsmmigrate**, **dsmrecall**, **dsmrecalld**, or the **dsmreconcile** command with the **preptwo** option.

hsmgroupedmigrate

The **hsmgroupedmigrate** option specifies whether the **dsmmigrate** command migrates more than one file with each transaction.

The transfer of multiple files per transaction can improve performance when migrating many small or medium size files (up to 100 MB).

The default action is to migrate each file in a single transaction and make a stub file when the transaction is successful.

When this option is set to YES the HSM transaction grouping is activated. A group of files is migrated in each transaction with the IBM Spectrum Protect server when either the transaction byte limit or the transaction group limit is reached. The transaction byte limit has a default of 25 MB, and can be set with the **txnbytelimit** option. The transaction group limit is specified with the **txngroupmax** option. The **txngroupmax** option is set on the server in the options file or on the node definition on the server. If a file with a different management class setting is added to the group, the group is migrated immediately.

If a file with a different file system root is added to the group, the group is migrated immediately.

The **hsmgroupedmigrate** option works in LAN-free configurations and when data is transferred over LAN.

Options File

Place this option in the client options file `dsm.opt`.

Syntax



Parameters

NO The **dsmmigrate** command migrates one file per transaction with the IBM Spectrum Protect server. This is the default.

YES

The **dsmmigrate** command migrates a group of files in each transaction with the IBM Spectrum Protect server.

hsmlogeventflags

The **hsmlogeventflags** option specifies which kinds of events are added to the HSM log file.

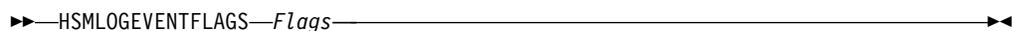
You must use an authorized user ID to use this option.

If this option is not specified, or if no flags are included in the specification, no HSM log entries are made.

Options File

Place this option in the client system options file `dsm.sys` within a server stanza.

Syntax



Parameters

Flags

Indicate one or more of the following values:

FILE

File events are logged. File events include file migration and recall processing.

FS File system events are logged. File system events include threshold migration, reconciliation, and HSM-related file system configuration changes.

hsmlogmax

The **hsmlogmax** option specifies the maximum size of the HSM log, in megabytes.

You must use an authorized user ID to use this option.

If you use the **hsmlogretention** option instead of the **hsmlogmax** option, all existing log entries are retained and the log is pruned according to the new **hsmlogretention** option age criteria.

If you use the **hsmlogmax** option instead of the **hsmlogretention** option, all records in the existing log are copied to the pruned log file, `dsmhsm.pru`. Then the existing log is emptied, and logging begins under the **hsmlogmax** option size criteria.

If you change the value of the **hsmlogmax** option, the existing log is extended or shortened to accommodate the new size. If the value is reduced, the oldest entries are deleted to reduce the file to the new size.

Restriction: You cannot specify a nonzero **hsmlogmax** option value and activate the **hsmlogretention** option.

Options File

Place this option in the client system options file `dsm.sys` within a server stanza.

Syntax

▶▶—HSMLOGMAX—*size*—————▶▶

Parameters

size

Specifies the maximum size, in megabytes, for the log file. The range of values is 0 - 2047; the default is 0, which specifies that the log file has no size limit.

Related reference:

“**hsmlogretention**” on page 100

hsmlogname

The **hsmlogname** option specifies the HSM log file name.

This option specifies the fully qualified path and file name of the HSM log file. If this option is not specified, the default log file is used. The default log file is `dsmhsm.log` and is in the same directory as the `dsmerror.log` file.

You must use an authorized user ID to use this option.

The value for this option overrides the **DSM_LOG** environment variable.

Options File

Place this option in the client system options file `dsm.sys` within a server stanza.

Syntax

▶—HSMLOGName—*filespec*—▶

Parameters

filespec

The fully qualified path and file name of the HSM log file. If any part of the path you specify does not exist, the space management client creates it.

The log file path cannot be a symbolic link.

hsmlogretention

The **hsmlogretention** option specifies how many days to maintain HSM log entries before pruning, and whether to save the pruned entries.

The HSM log is pruned when the first entry is written to the log after a space management client command or daemon is started. Some daemons run continuously. Stop the daemons and start them again or activate one of the HSM commands to prune the HSM log.

You must use an authorized user ID to use this option.

If you use the **hsmlogmax** option instead of the **hsmlogretention** option, all records in the existing log are copied to the pruned log file, `dsmhsm.pru`. Then the existing log is emptied, and logging begins under the new size criteria.

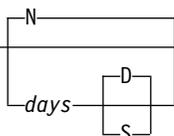
If you use the **hsmlogretention** option instead of the **hsmlogmax** option, all existing log entries are retained and the log is pruned according to the new **hsmlogretention** option age criteria.

Restriction: You cannot specify the **hsmlogretention** option and a nonzero value for the **hsmlogmax** option.

Options File

Place this option in the client system options file `dsm.sys` within a server stanza.

Syntax

▶—HSMLOGRetention—▶

Parameters

N Do not prune the log. The HSM log grows indefinitely. This is the default.

days

The number of days to keep log file entries before pruning the log. The range of values is 0 - 9999.

D or S

Specifies whether to save the pruned entries. You must also specify the *days* parameter. Enter a space or comma to separate the *days* parameter from the *D* or *S* parameter.

D Discard the HSM log entries when you prune the log. This is the default.

Example task: Prune the log every 7 days. Discard the pruned log entries.

Example command: `hsmlogr 7 d`

S Save the HSM log entries when you prune the log. The pruned entries are copied from the HSM log to the `dsmhsm.pru` file that is in the same directory as the HSM log.

Example task: Prune the log every 100 days. Save the pruned log entries.

Example command: `hsmlogr 100,s`

Related reference:

“`hsmlogmax`” on page 99

Related information:

Chapter 12, “HSM client command reference,” on page 113

hsmlogsampleinterval

The **hsmlogsampleinterval** option specifies how often to gather file system data for the HSM log file.

You must be an authorized user to use the **hsmlogsampleinterval** option.

Options File

Place this option in the client system options file `dsm.sys` within a server stanza.

Syntax

▶▶—`HSMLOGSampleinterval—interval`—————▶▶

Parameters

interval

Specifies the number of seconds between each gathering of file system data.

The range of values is 0 - 9999999 and the default value is 3600 (one hour). If 0 is specified, no file system data is gathered.

hsmmaxrecalltapedrives

The **hsmmaxrecalltapedrives** option controls how many tape drives participate in tape-optimized recall processing.

By limiting the available tape drives for tape optimized recall processing, you can prevent all drives from being blocked. In this way, some drives remain accessible for transparent recall processing and migration.

If the value of the **hsmmaxrecalltapedrives** option is greater than the number of tape drives available, a recall thread waits for the next available tape drive. If you

plan to run tape-optimized recalls processes in parallel, the number can be decreased accordingly for parallel processing without running out of tape drives.

If the **maxnummp** option is defined on the IBM Spectrum Protect server, you can set the **hsmmaxrecalltapedrives** option to a number that is less than or equal to the **maxnummp** option. You can reserve tape drives for other operations if you set the number lower than the **maxnummp** option and the maximum number of tape drives.

Options File

Place this option at the beginning of the `dsm.sys` file before any server stanzas.

Syntax

▶▶—HSMMAXRECALLTAPEDRIVES—*value*————▶▶

Parameters

value

Specifies the number of tape drives available for a tape-optimized recall. The range of values is 1 -10. The default is 5.

hsmmigzeroblockfiles

The **hsmmigzeroblockfiles** option specifies whether the space management client migrates files that allocate zero data blocks in the file system.

Some files are so small that the metadata and the file content fit completely in the inode. The file system allocates no data blocks for such small files. Migrating such small files does not save space on the file system. By default, the space management client migrates such small files when the files meet migration criteria. You can set HSMMIGZEROBLOCKFILES NO if you do not want to migrate such small files.

Options File

Place this option at the beginning of the `dsm.sys` file before any server stanzas.

Syntax

▶▶—HSMMIGZEROBLOCKFILES—

YES
NO

————▶▶

Parameters

YES

Files that allocate zero data blocks in the file system, but have a logical size greater than zero, can be migrated. This is the default.

NO Files that allocate zero data blocks in the file system cannot be migrated.

hsmmultiserver

The **hsmmultiserver** option enables and disables migration and backup in an environment of multiple IBM Spectrum Protect servers.

Supported Clients

You can enable migration and backup from a single file system to multiple IBM Spectrum Protect servers. You must set **hsmmultiserver=YES** for each space management client in the GPFS cluster.

If you set **hsmmultiserver=NO**, all of the command options for the multiple-server environment are disabled. When you set **hsmmultiserver=NO**, you can recall files only from the IBM Spectrum Protect migration server that is configured in the `dsm.sys` options file. If a file was migrated to a server other than the server that is configured in the `dsm.sys` options file, you cannot recall the file.

The option is available only for AIX GPFS clients and Linux GPFS clients.

Options File

Place this option in the client options file `dsm.opt`.

Syntax

►►—HSMMULTIServer—NO
YES—►►

Parameters

NO All of the command options for the multiple-server environment are disabled. All HSM commands have standard input and output. This is the default.

YES

All of the command options for the multiple-server environment are enabled. All HSM commands that display a list add an additional column that shows the IBM Spectrum Protect server name.

Related information:

Chapter 12, “HSM client command reference,” on page 113

inlxc1

Use the **inlxc1** option to define the filename and path of your include-exclude options file.

Syntax

►►—inlxc1—*filespec*—►►

Parameters

filespec

Specifies the path and file name for your include-exclude file.

Related concepts:

“Include-exclude file options” on page 26

Related tasks:

“Creating an include-exclude list” on page 28

maxcandprocs

The **maxcandprocs** option specifies the number of parallel threads in the scout daemons that can scan for file systems.

Tip: When automatic migration is driven by GPFS policy, the **dsmscoutd** daemon is disabled, and the **maxcandprocs** option has no effect.

Options File

Place this option at the beginning of the `dsm.sys` file before any server stanzas.

Syntax

▶▶—MAXCANDProcs—*number*————▶▶

Parameters

number

Specifies the maximum number parallel threads in the scout daemons that can scan file systems. The range of values is 2 - 20. The default is 5.

maxmigrators

The **maxmigrators** option specifies the maximum number of parallel migration sessions that can run in parallel for each file system. Ensure that you have sufficient resources on the server for parallel migration to occur.

Tip: When automatic migration is driven by GPFS policy, the **maxmigrators** option has no effect.

Do not set the **maxmigrators** option higher than the number of parallel sessions that the server can use to store data.

This option can be set by the IBM Spectrum Protect server.

Options File

Place this option at the beginning of the `dsm.sys` file before any server stanzas.

Syntax

▶▶—MAXMIGRators—*number*————▶▶

Parameters

number

Specifies the maximum number of parallel migration sessions that you can set. The range is 1 - 20. The default is 5. If this option is changed from the default,

make a corresponding increase in the IBM Spectrum Protect server configuration to update the HSM node **MAXNUMMP** value.

maxrecalldaemons

The **maxrecalldaemons** option specifies the maximum number of recall daemons that you can run at one time to recall files for the client node. If the number of recall daemons that are running at one time is close to the maximum number, increase the value.

For example, if you use an application that opens many migrated files at once, the application can use all available recall daemons. Another process cannot access a migrated file until a recall daemon is available.

Note: The GPFS configuration option **dmapiWorkerThreads** can limit the maximum concurrent recall tasks to less than the value of **maxrecalldaemons**.

Options File

Place this option at the beginning of the `dsm.sys` file before any server stanzas.

Syntax

▶▶—MAXRECALLdaemons—*number*—————▶▶

Parameters

number

Specifies the maximum number of recall daemons that can run in parallel to recall files for the client node. The range is 2 - 99. The default is 20.

maxthresholdproc

The **maxthresholdproc** option specifies the maximum number of threshold migration processes that the space management client can start at one time. When a file system runs out of space, the HSM client does not verify the maximum number of threshold migration processes that currently are running. It starts threshold migration as part of the demand migration process *regardless* of the number of threshold migration processes in progress.

Tip: When automatic migration is driven by GPFS policy, the **maxthresholdproc** option has no effect.

Options File

Place this option at the beginning of the `dsm.sys` file before any server stanzas.

Syntax

▶▶—MAXThresholdproc—*number*—————▶▶

Parameters

number

Specifies the maximum number of automatic threshold migration processes that the space management client can start at one time. The range of values is 1 - 99. The default is 3.

migfileexpiration

The **migfileexpiration** option specifies the number of days that copies of migrated or premigrated files remain on the server after they are modified on your local file system or are deleted from your local file system.

Options File

Place this option at the beginning of the `dsm.sys` file before any server stanzas.

Syntax

►►—MIGFileexpiration—*days*—————►►

Parameters

days

Specifies the number of days a copy of a migrated or premigrated file remains in storage after it is modified on your local file system, or deleted from your local file system. The range of values is 0 - 9999. The default is 7 days.

Note: If you specify a value of 0, an obsolete copy of a migrated or premigrated file is deleted from the server during the next reconciliation run. If you delete a file from the local file system and reconcile runs with the `migfileexpiration` value as 0, the file can not be recreated by the **dsmmigundelete** process.

migrateserver

The **migrateserver** option specifies the name of the server to which you want to migrate files from your client node. Specify one migration server for each client node.

If you do not specify a server with the **migrateserver** option, your files migrate to the server that you specify with the **defaultserver** option. If you do not specify a server with either of these options, your files migrate to the server that you identify in the first stanza of your `dsm.sys` file.

You can override this option with the following command: `dsmmigfs upd /FS -Server=servername`. Replace `servername` with the name of your server.

The value of **migrateserver** in `dsm.sys` overrides **migrateserver** in `dsm.opt`.

After your files migrate to the server that you specified, do not specify a different migration server unless your administrator transfers your migrated files from the specified server to another. Otherwise, the server cannot locate your migrated files until you specify the server to which your files were originally migrated.

Options File

Place this option at the beginning of the `dsm.sys` file before any server stanzas.

Syntax

►►—MIGRateserver—*servername*—————►►

Parameters

servername

Specifies the name of the server to which you want to migrate files from your client node. Your `dsm.sys` file must contain a stanza beginning with the **servername** option and it must contain the required communication options for the server that you specify with the **migrateserver** option.

Related reference:

“**defaultserver**” on page 92

minmigfilesize

The **minmigfilesize** option specifies the minimum file size for a file to be eligible for migration.

This option applies to all of the space-managed file systems for which you have not specified a file system specific value for **minmigfilesize** using the **dsmmigfs add** or **dsmmigfs add** commands. See “**dsmmigfs add** and **update**” on page 128 for details.

Options File

Place this option at the beginning of the `dsm.sys` file before any server stanzas.

Syntax

►►—MINMIGfilesize—*fileSize*—————►►

Parameters

fileSize

Specifies the minimum file size, in bytes, for a file to be eligible for migration. The range of values is 0 through 2147483647. The default is 0.

On GPFS file systems, if you specify the default, the space management client uses the current file system stub size as the minimum size for files that can be migrated. For other (non-GPFS) file systems, the space management client uses the file system block or fragment size or stub size as the minimum size for files that can be migrated, whichever is larger.

If you specify a non-zero value, for GPFS file systems, the value must be greater than the stub size. For other (non-GPFS) file systems, it must be greater than both the file system block or fragment size or stub size. Otherwise, the value is ignored.

minrecalldaemons

The **minrecalldaemons** option specifies the minimum number of recall daemons that can run in parallel to recall files for the client node.

Options File

Place this option at the beginning of the `dsm.sys` file before any server stanzas.

Syntax

▶▶—MINRecalldaemons—*number*—————▶▶

Parameters

number

Specifies the minimum number of recall daemons that can run in parallel. The range of values is 1 - 99. The default is 3.

reconcileinterval

The **reconcileinterval** option specifies how often the space monitor daemon reconciles your file systems. Depending on the **checkfororphans** option, the reconciliation either expires or deletes obsolete objects on the server and updates the status file, or checks for orphan stub files and makes metadata updates.

Tip: When **hsmdisableautomigdaemons=YES**, the **reconcileinterval** option has no effect.

Options File

Place this option at the beginning of the `dsm.sys` file before any server stanzas.

Syntax

▶▶—RECOncileinterval—*interval*—————▶▶

Parameters

interval

Specifies the number of hours that must elapse between each successive time your file systems are automatically reconciled on your workstation. If you specify a value of 0, your file systems are not reconciled automatically. The range of values is 0 - 9999. The default is 24.

restoremigstate

The **restoremigstate** option specifies whether you want to restore or retrieve stub files or backup-archive versions of migrated files during a restore-retrieve operation. Use this option with the backup-archive client **restore** and **retrieve** commands.

You can restore or retrieve a stub file for a migrated file only when the following is true:

- The file exists in the migration storage pool
- The file is backed up or archived and migrated to the same server

When the number of days elapse that you specified with the **migfileexpiration** option, the migrated file is removed from storage.

If you specify **restoremigstate YES**, and if the migrated file is not expired, the file is restored or retrieved to a stub file, regardless of whether it is marked for expiration.

On a space-managed file system, you can restore a stub file only during standard-query restore processing. During no-query restore processing to a space-managed file system, you cannot restore stub files.

The **restoremigstate** option restores a file if it is backed up after migration. If the file is backed up before migration, you cannot restore a stub file because a server stub file copy does not exist.

Files with access control lists (ACLs) are restored in a resident state regardless of the setting for **restoremigstate**. The restore operation affects files that are restored to a GPFS file system with more storage pools than the default system pool.

If you restore a file from the IBM Spectrum Protect target server, and the file system is managed by the space management client, you must not restore the file as a stub file. You must restore the complete file. Use the **restoremigstate=no** option to restore the complete file. If you restore the file as a stub from the target server, the following consequences can occur:

- You cannot recall the file from the IBM Spectrum Protect source server by using the space management client.
- A space management client reconciliation process that runs against the IBM Spectrum Protect source server expires the file. If the file is expired by a reconciliation process, you can restore the complete file with the backup-archive client and the **restoremigstate=no** option.

Tip: The **restoremigstate** option does not support hard linked files. If you want to restore or retrieve a stub file for a hard linked file, delete all files from your local file system that are hard linked together. When one file in a set of hard linked files is migrated, all of the hard linked files in the set become stub files. When you enter the **restore** command with the **restoremigstate** option to restore a stub file for a hard linked file, the stub file has the same name as the file that was originally migrated. Stub files are not restored for any other files that previously were in the hard linked set of files.

Options File

Place this option in the **dsm.opt** or **dsm.sys** file.

Syntax



Parameters

YES

Restores or retrieves migrated files to stub files on your local file system during a restore or retrieve operation. The files remain migrated. This is the default.

Note: A stub file that is created during a restore or retrieve operation contains the information that is necessary to recall the migrated file from storage. It does not contain any leading bytes of data from the file. Any recall mode that was set previously for the migrated file (for example, streaming or partial file recall) is not stored in the stub file. The recall mode is set to normal for all files that are restored or retrieved to stub files.

NO Restores or retrieves backup-archive versions of migrated files to your local file system during a restore or retrieve operation. The files become resident.

Command line

This option is valid on the command line.

Related tasks:

“Restoring migrated files” on page 65

skipmigrated

The **skipmigrated** option specifies whether the IBM Spectrum Protect backup-archive client process migrated files during a backup or archive operation. Use this option with the backup-archive client **archive**, **backup image**, **incremental**, and **selective** commands.

When the **skipmigrated** option is set to yes, the backup-archive client does not back up or archive any stub files. When the **skipmigrated** option is set to no, the backup-archive client can process stub files during some operations.

Options File

Place this option in the `dsm.opt` file.

Syntax



Parameters

NO The backup-archive client does not check the migration status of files during a backup or archive operation. This is the default.

YES

The backup-archive client checks the migration status of files during a backup or archive operation. The backup-archive client does not back up or archive stub files.

Command line

This option is valid on the command line.

Related tasks:

“Backing up migrated files” on page 63

Chapter 12. HSM client command reference

You can use commands to do all HSM tasks. You must use the correct syntax and format when you enter commands.

When you issue commands and options, follow these rules:

- Do not precede HSM commands with **dsmc**. Each space management client command is a separately executable command.
- Enter the complete command name in lowercase letters. You cannot use uppercase letters or an abbreviation for a command name.
- Do not stop any HSM process with the command `kill -9`.
- Use the following wildcard characters in file, directory, or file system specifications. The shell in which you are running matches and expands wildcard characters.
 - * Matches zero or more characters
 - ? Matches any single character
- Enter the characters in a command in a continuous string without pressing the **Return** key. You can enter as many as 256 characters on the command line.

Related concepts:

Chapter 3, “Configuring the space management client,” on page 19

Standard option formats

You can specify options on some commands. You must use a standard format to specify all options.

Follow these guidelines when you use an option:

- Uppercase letters in each option description indicate the minimum abbreviation of a an option. Type the complete option name or an abbreviation of the name.
- Issue options in any combination of uppercase and lowercase letters. Options are not case-sensitive.
- Precede each option with a hyphen (-). For example:

```
dsmmigquery -mgmtclass -detail /home
```
- Separate each option with a blank space.
- Issue more than one option in a command in any order before or after a file, directory, or file system specification.
- If the option defines a value, separate the option name from the value with an equal sign (=). For example:

```
dsmmigfs update -ht=90 /home
```

The table gives examples of options that do not define a value.

Table 15. Option format examples: options without values

Command
<code>dsmmigrate -recursive -detail /home/user1/file1</code>
<code>dsmmigrate -rec -det /home/user1/file1</code>
<code>dsmmigrate -r -d /home/user1/file1</code>
<code>dsmmigrate /home/user1/file1 -r -d</code>

The table gives examples of options when a value is defined.

Table 16. Option format examples: options with values

Command
<code>dsmmigfs update -hthreshold=90 -lthreshold=20 /home</code>
<code>dsmmigfs update -hthresh=90 -lthreshold=20 /home</code>
<code>dsmmigfs update -ht=90 -l=20 /home</code>
<code>dsmmigfs update /home -ht=90 -l=20</code>

In version 6.4 and later, the **optionformat** option is not supported. Short option format is not supported. You must specify all options with the standard option format.

Help for commands

Use the **help** option to display help for each command-line command or use the **dsmmighelp** command.

You can display online help for HSM commands in either of the following ways:

- Issue the help option with any command. For example:

```
dsmmigrate -help
dsmmigrate -h
```
- Issue the **dsmmighelp** command. A list of help topics displays from which you can select general help information for commands, help for a specific command, or help for a message.

Proper display of the help text requires a usable display width of 72 characters. A display width that is fewer than 72 characters causes sentences that are 72 characters wide to wrap to the next line. Wrapping can cause the displayed help text to begin somewhere within the section rather than at the beginning. The skipped lines can be viewed by using the terminal's scrolling function to move up.

Displaying file and file system information

There is a set of HSM commands you can use to display space management information about your file systems, files, and directories.

Table 17. HSM commands to display file and file system information

Command	Description
dsmdf	Displays space usage information for a file system. For example, to display space usage information for the /home file system, issue the following command: <pre>dsmdf /home</pre> See “ dsmdf ” on page 123 for more information about this command.
dsm1s	Lists files in a directory and displays file conditions. For example, to display information about all files in the /home/user1 directory, issue the following command: <pre>dsm1s /home/user1/*</pre> See “ dsm1s ” on page 126 for more information about this command.

Table 17. HSM commands to display file and file system information (continued)

Command	Description
dsmdu	<p>Displays space usage information for files and directories. For example, to display space usage information for each file in the /home/user/proj1 directory and in all of its subdirectories, issue the following command:</p> <pre>dsmdu -Allfiles /home/user1/proj1</pre> <p>See “dsmdu” on page 124 for more information about this command.</p>
dsmmigfs query	<p>Displays the current space management settings for a file system. For example, to display the space management settings for the /home file system, issue the following command:</p> <pre>dsmmigfs query /home</pre> <p>See “dsmmigfs query” on page 139 for more information about this command.</p>
dsmmigundelete	<p>Recreates deleted stub files for migrated files and creates stub files for premigrated files if a corresponding original file does not exist on your local file system. The file then becomes a migrated file. For example, to recreate stub files for migrated files in the /home file system that are not marked for expiration (reconciliation was not run since the files were deleted) issue the command:</p> <pre>dsmmigundelete /home</pre> <p>See “dsmmigundelete” on page 153 for more information about this command.</p>

Client return codes

The space management client command-line interface exits with return codes that accurately reflect the success or failure of the operation.

Scripts, batch files, and other automation facilities can use the return code from the command-line interface. For operations that use the IBM Spectrum Protect scheduler, the return codes are shown in the output of the **QUERY EVENT** administrative command. For cases where the return code is not 0, you can examine the `dsmerror.log` file. For scheduled events, you can examine the `dsmsched.log` file.

Return codes have the following meanings:

Table 18. An explanation of client return codes

Code	Explanation
0	All operations completed successfully.
4	<p>The operation completed successfully, but some files were not processed. There were no other errors or warnings. This return code is common. In most cases, files are not processed for the following reasons:</p> <ul style="list-style-type: none"> • The file satisfies an entry in an exclude list. Excluded files generate log entries only during selective backups. • The file was in use by another application and could not be accessed by the client. • The file changed during the operation to an extent prohibited by the copy serialization attribute.

Table 18. An explanation of client return codes (continued)

Code	Explanation
8	The operation completed with at least one warning message. Review the <code>dsmerror.log</code> file to determine what warning messages were issued and to assess their effect on the operation.
12	The operation completed with at least one error message (except for error messages for skipped files). For scheduled events, the status is Failed. Review the <code>dsmerror.log</code> file to determine what error messages were issued and to assess their effect on the operation. Generally, this return code means that the error was severe enough to prevent the successful completion of the operation. For example, an error that prevents an entire file system from being processed yields return code 12.

The return code for a client macro is the highest return code issued among the individual commands that comprise the macro. For example, suppose that a macro consists of these commands:

```
selective "/home/devel/*" -subdir=yes
incremental "/home/devel/TestDriver/*" -subdir=yes
archive "/home/plan/proj1/*" -subdir=yes
```

If the first command completed with return code 0; the second command completed with return code 8; and the third command completed with return code 4, the return code for the macro is 8.

Related tasks:

 [Configuring backup-archive clients](#)

Related information:

 [Copy serialization attribute](#)

 [QUERY EVENT administrative command](#)

HSM command summary

Table 19 provides an alphabetical list of the space management client commands, a brief description of each command, and the command page number.

Table 19. HSM command summary

Command and location	Description
dmkilld	Valid only on GPFS file systems. Stops the master recall daemon and all of its child processes, and interrupts all active recalls. See “ dmkilld ” on page 118.
dsmattr	Valid only on GPFS file systems. Sets or displays the recall mode for a migrated file. See “ dsmattr ” on page 118.
dsmautomig	Starts parallel migration sessions for a file system. See “ dsmautomig ” on page 121.
dsmdf	Displays space usage information for a file system. See “ dsmdf ” on page 123.
dsmdu	Displays space usage information for files and directories. See “ dsmdu ” on page 124.
dsm ls	Lists files in a directory and displays file state. See “ dsm ls ” on page 126.
dsmmigfs add, update	Adds space management to a file system, or updates space management attributes for a file system. See “ dsmmigfs add and update ” on page 128.

Table 19. HSM command summary (continued)

Command and location	Description
dsmmigfs addmultiserver, querymultiserver, removemultiserver	Adds space management to a file system, or updates space management attributes for a file system. See “ dsmmigfs addmultiserver, querymultiserver, and removemultiserver ” on page 134.
dsmmigfs deactivate, reactivate, remove	Deactivates or reactivates space management for a file system, or removes space management from a file system. See “ dsmmigfs deactivate, reactivate, and remove ” on page 137.
dsmmigfs enablefailover, disablefailover	Manages recovery from partial system failure (GPFS only). See “ dsmmigfs enablefailover, and disablefailover ” on page 143.
dsmmigfs query	Displays current space management settings for a file system. See “ dsmmigfs query ” on page 139.
dsmmigfs globaldeactivate, globalreactivate	Deactivates or reactivates space management for a space-managed client node. See “ dsmmigfs globaldeactivate and globalreactivate ” on page 138.
dsmmigfs rollback	Valid only on GPFS file systems. Transfers the HSM management of a file system to the preferred node if the node is different from the current owner node. See “ dsmmigfs rollback ” on page 142.
dsmmigfs stop, start, restart	Starts or starts the HSM daemons. See “ dsmmigfs stop, start, and restart ” on page 144.
dsmmigfs takeover	Transfers the HSM management of a file system to a space management client node within the same local GPFS cluster. See “ dsmmigfs takeover ” on page 146.
dsmmighelp	Displays online help for commands. See “ dsmmighelp ” on page 147.
dsmmigquery	Displays space management information. There are many HSM client and backup-archive client shared options available using the dsmmigquery -o command. See “HSM and backup-archive client dsmmigquery command shared options” on page 149 for a list of these shared options. See “ dsmmigquery ” on page 147.
dsmmigrate	Moves selected files from your local file system to IBM Spectrum Protect storage. See “ dsmmigrate ” on page 149.
dsmmigundelete	Re-creates deleted stub files. See “ dsmmigundelete ” on page 153.
dsmmonitord	Starts the space monitor daemon. See “ dsmmonitord ” on page 155.
dsmq	Displays information, including recall IDs, for all files that are currently queued for recall. See “ dsmq ” on page 156.
dsmrecall	Moves selected files from storage to your local file system. See “ dsmrecall ” on page 157.
dsmrecalld	Starts the recall daemon. See “ dsmrecalld ” on page 160.
dsmreconcile	Synchronizes the client and server. See “ dsmreconcile ” on page 161.
dsmrm	Removes a recall process from the recall queue. See “ dsmrm ” on page 163.
dsmscoutd	Starts, stops and restarts the scout daemon and shows file system information. See “ dsmscoutd ” on page 164.
dsmsetpw	Changes the IBM Spectrum Protect password for your client node. See “ dsmsetpw ” on page 165.

Table 19. HSM command summary (continued)

Command and location	Description
dsmwatchd	Valid only on GPFS file systems. The dsmwatchd command starts the watch daemon. The watch daemon checks the status of the recall daemon, the monitor daemon, and the scout daemon. If any of these daemons end or become corrupted, the watch daemon automatically recovers the failed daemon. See “ dsmwatchd ” on page 166.

dmkilled

The **dmkilled** command stops the master recall daemon and all of its children and interrupts all active recalls.

This command is valid only on GPFS file systems.

Syntax

```

▶▶—DMKILLD—┬──options──┴──▶▶

```

Parameters

options

—help

Displays syntax and options for the command. Do not specify other options when you specify the **help** option.

dsmattr

The **dsmattr** command sets or changes the recall mode for one or more migrated files. The recall mode determines how the space management client recalls a migrated file when you access it.

You cannot set a recall mode for a resident or a premigrated file. The recall mode that you set for a migrated file remains associated with that file only while the file remains migrated.

If you do not specify any parameters, the command displays the current recall mode, stub size, preview size, file name, and whether the **readstartsrecall** option is set.

Syntax

```

▶▶—DSMATTR—┬──options──┴──┬──filespec──┴──▶▶
                                     └──Filelist—==file_path┘

```

Parameters

options

-ERRORLOGName=*file_path*

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file `dsm.sys` or `dsm.opt`, or as specified by the **DSM_LOG** environment variable. If any part of the path you specify does not exist, the space management client creates it.

-Help

Displays syntax and options for the command. Do not specify other options when you specify the **help** option.

-PREViewsize=*n*

Specifies the preview size of the stub file. Valid values are 0 - 1,073,741,824, and must not be greater than the stub size.

When the attributes are set and the preview size is defined, a recall operation begins only when an application reads the stub file beyond the preview section of the stub file. If an operation reads only the preview section of the stub file, the file is not recalled.

-REAdstartsrecall=No|Yes

Specifies whether a recall operation starts when an application reads the stub file.

This option applies only if streaming recall mode is set.

If the **previewsize** option is set for the stub file, a recall operation starts only when an application reads beyond the preview section of the stub file.

The default value is No.

Regardless of the value of the **readstartsrecall** option, the file is recalled when either of the following conditions are true:

- The recall daemon determines that the stub file does not contain all of the data that is required by a read operation.
- An operation modifies the file.

-RECAllmode=*value*

Sets a recall mode for one or more migrated files. If you do not specify the **RECAllmode** option, the space management client displays the current recall mode for the files that you specify. You can specify the following values:

Value	Description
Normal	Recalls the migrated file to its originating file system. Normal is the default. If the file is not modified, it becomes a premigrated file. If the file is modified, it becomes a resident file.
Partialrecall GPFS systems only)	Specifies that the file is recalled by using partial file recall, regardless of its size.
Streaming	Specify an asynchronous recall of migrated files. The recalled portion of the file can be accessed before the complete file is recalled. This parameter is valid for read-only operations on the file.

-RECUrsive

Sets or displays the recall mode for migrated files in the directory and subdirectories that you specify.

filespec

Note: This parameter is required only when you set a new recall mode. If you do not use the **RECALLmode** option, and you do not specify a path and a file name, the current recall mode displays for all files in the current directory.

The path and file name of the file for which you want to set a new recall mode, or display the current recall mode. You can specify a single file, a group of files, or a directory. If you specify a directory, the space management client sets or displays that recall mode for each migrated file in the directory.

You can use wildcard characters to specify a group of files with similar names. You can issue more than one file specification in a command. If you enter several file specifications, separate each specification with one or more blank spaces.

-Filelist=*file_path*

Specifies the path of a file that contains a list of files to be processed by this command.

The entries in the list file must adhere to the following rules:

- Each entry is a fully qualified path to a file or a relative path to a file.
- No entry is a directory object.
- Each entry is on a separate line.

The space management client ignores any entry that does not adhere to these rules.

The following rules apply to a file list for the space management client. These rules differ from the rules for backup-archive client file lists.

- An entry can contain unprintable characters, but cannot contain a carriage return.
- Wildcard characters are allowed.
- A file path that contains blank spaces must be enclosed in quotation marks.
- An entry that begins with a quotation mark and ends with a quotation mark is tolerated. The space management client assumes that the quotation marks are not needed and omits them when the space management client processes the entry. If beginning and ending quotation marks are required to identify the file, add double quotation marks before the beginning and after the end of the entry.

The following example shows a list of files within a list file:

```
/home/dir/file1
"/fs1/dir2/file3"
"/fs2/my files/file4"
../tivoli/'file1'
"'fs3'/dir3/'file.txt'"
fs4/dir/a"file".txt
'/fs4/dir/file.txt'
/fs5/dir/file*with?wildcards.txt
```

If the list file name that you specify with the **filelist** option does not exist, the command fails.

Examples

Task	Command
Change the recall mode to partial file recall for all migrated files in the /home/user2 directory and all of its subdirectories.	<code>dsmattr -recall=partialrecall -Recursive /home/user2</code>

Task	Command
Enable an asynchronous recall of migrated files in the <code>/home/user2/</code> directory.	<code>dsmatrr -recall=streaming /home/user2/</code>
On the <code>/myfs</code> file system, stub files are 32,768 bytes. For the migrated files that are listed in <code>files.lst</code> , set the preview size to 4,096 bytes. If an application reads more than the preview section of any of these stub files, start the recall process. (Streaming recall mode is a requirement to start the recall after an application reads the preview section).	<code>dsmatrr -preview=4096 -readstartsrecall=yes -recall=streaming -filelist=/myfs/files.lst</code>
Display the recall attributes that are assigned to all files in the current directory.	<code>dsmatrr</code>
Display all files in the file list that is named <code>/tmp/filelist</code> .	<code>dsmatrr -filelist=/tmp/filelist</code>

Related concepts:

Chapter 7, “Recalling migrated files,” on page 71

dsmautomig

The **dsmautomig** command starts parallel migration sessions to the IBM Spectrum Protect server, migrating more than one file at a time.

You must have root user authority to use this command.

The **dsmautomig** command checks:

- If a migration candidate requires a current backup version on the TDP for IBM Lotus Notes® IBM Spectrum Protect server.
- If a current backup version exists.

If the `LANG` environment variable is set to `C`, `POSIX`, only characters with ASCII codes less than 128 are valid. The HSM client skips files that have file names that contain invalid characters. If you are using a single-byte character set (SBCS) such as English as your language environment, all file names are valid and are migrated by the space management client.

Multi-byte characters are interpreted as a set of single bytes all containing valid characters. If you are using multi-byte character sets (MBCS) as your language environment, the HSM client migrates file names that consist of valid characters in the current environment. For example, a file name consisting of Japanese characters might contain invalid multi-byte characters if the current language environment is a Chinese character set. File names containing invalid multi-byte characters are not migrated or recalled. If such files are found during migrate or recall no information is printed. The HSM daemons must run in the `en_US` language locale to work properly.

Specify the number of parallel migration sessions with the **maxmigrators** option in your `dsm.sys` file. Verify that sufficient resources are available on the IBM Spectrum Protect server for parallel migration. Do not set the **maxmigrators** option higher than the number of sessions that the IBM Spectrum Protect server can use to store data. Start threshold migration manually to reduce space usage on your file system before it reaches the high threshold that you set.

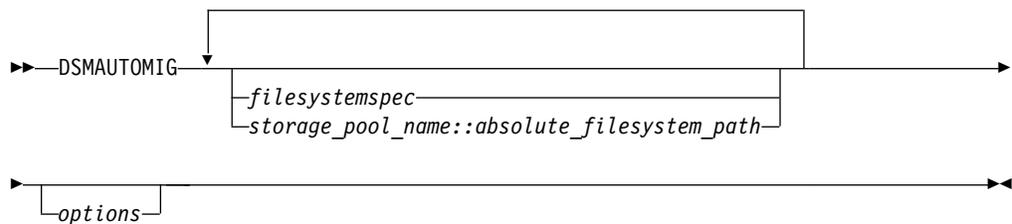
Hidden directories and files are included in automatic migration. Files can be excluded from automatic migration by adding the hidden directories or files to the exclude list in the `dsm.opt` file.

The scout daemon (`dsmscoutd`) should be running if you start the **dsmautomig** command manually. Otherwise, the **dsmautomig** command might not be able to complete the migration if it runs out of candidates from the candidates list.

The **dsmautomig** command must be found with the `PATH` variable, or the **dsmonitord** daemon cannot complete threshold migration.

Note: The HSM client does not migrate contents of symbolic links.

Syntax



Parameters

filesystemspec

Specifies the name of the file system for which you want to run threshold migration. The default is all file systems for which space management is active. You can specify more than one file system name, and you can use wildcard characters within a file system name. If you specify more than one file system name, separate each name with one or more blank spaces.

storage_pool_name::absolute_filesystem_path

Specifies the storage pools that are in the *absolute_filesystem_path* path that are migrated automatically.

options

-Detail

Displays information about migrated files.

-ERRORLOGName=*file_path*

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file `dsm.sys` or `dsm.opt`, or as specified by the **DSM_LOG** environment variable. If any part of the path you specify does not exist, the space management client creates it.

-Help

Displays syntax and options for the command. Do not specify other options when you specify the **help** option.

Examples

Task	Command
Start threshold migration for all storage pools on all file systems for which space management is active.	<code>dsmautomig</code>
Start threshold migration for all storage pools in the <code>/home</code> file system.	<code>dsmautomig /home</code>
Start threshold migration for all storage pools in the <code>/home</code> and <code>/test1</code> file systems.	<code>dsmautomig /home /test1</code>
Start threshold migration for the storage pools that are named <code>silver</code> and <code>gold</code> for the <code>/fs1</code> file systems.	<code>dsmautomig /silver::<fs1 code="" gold::<fs1<=""></fs1></code>
Start threshold migration for all storage pools in the <code>/fs2</code> file systems and for the storage pool that is named <code>gold</code> in the <code>/fs1</code> file systems.	<code>dsmautomig gold::<fs1 code="" fs2<=""></fs1></code>

dsmdf

The **dsmdf** command displays the information for one or more file systems, such as file system state, inode information, and space information.

Specifically, the **dsmdf** command displays information about:

- File system state: active (a), inactive (i), or global inactive (gi)
- Sum of the sizes of all migrated files
- Amount of space that is used on your local file system for premigrated files
- Number of inodes that are used for migrated or premigrated files
- Number of unused inodes on your local file system
- Amount of free space on your local file system

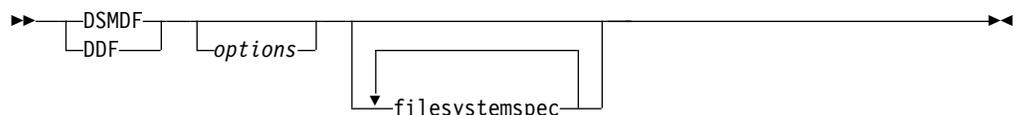
Only migration and recall processes dynamically update status information for your file systems. If any other process changes the state of a file, the change is not reflected in the information that the **dsmdf** command displays until reconciliation is run.

For GPFS file systems only:

Migrating newly created files less than five (5) minutes old might display incorrect results (resident size) when you use the **dsmdf** and **dsmdu** commands. The reason is because GPFS is not synchronized on all nodes when you migrate files. The last block of a file is not released from the disk although the file migrated successfully. This procedure can cause a deviation from an assumed disk usage if many small files are migrated and the block size is high.

Note: You can display only information about mounted file systems. If a file system is space managed but not mounted, it does not appear within the command output.

Syntax



Parameters

options

-Detail

Displays information about file systems with each value appearing on its own line. Values representing the amount of space is shown in kilobytes only.

-ERRORLOGName=*file_path*

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file `dsm.sys` or `dsm.opt`, or as specified by the **DSM_LOG** environment variable. If any part of the path you specify does not exist, the space management client creates it.

-Help

Displays syntax and options for the command. Do not specify other options when you specify the **help** option.

-Logname=*file_path*

Specifies the path and file name of the log file to be used by this command. The kinds of events that are logged to the file are specified with the **hsmlogeventflags** option. Error events are not logged to the file.

filesystemspec

The name of the file system for which you want to display information. The default is all file systems to which you added space management. You can specify more than one file system name, and you can use wildcard characters within a file system name. If you specify more than one file system name, separate each name with one or more blank spaces.

Examples

Task	Command
Display information for all file systems to which you added space management.	<code>dsmdf</code>
Display information for the /home file system.	<code>dsmdf /home</code>
Specify the log file to be used by the dsmdf command	<code>dsmdf -Logname=mylogfile</code> <code>dsmdf -detail -Logname=/tmp/dsmdflog /home</code> <code>dsmdf -d -L=mylogfile</code>

dsmdu

The **dsmdu** command displays space usage information for files and directories. For migrated files, the **dsmdu** command uses the actual size of the files that are stored in IBM Spectrum Protect storage to calculate space usage. In contrast, the **du** command (provided with your operating system) uses the size of the stub files that are stored in your local file system.

This command is valid only on GPFS file systems.

The **dsmdf** and **dsmdu** commands can display incorrect resident file size when you migrate newly-created files. This is because GPFS is not synchronized on all nodes when you migrate files. The last block of a file is not released from the disk

although the file migrated successfully. This can cause deviation from assumed disk usage if many small files are migrated and blocksize is high.

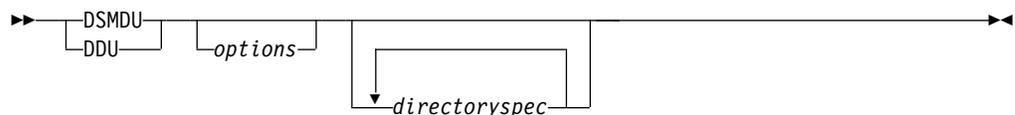
Attention: Running `dsmls` or `dsmdu` on remotely mounted GPFS file systems might show incorrect values. See Table 20 for an example of this situation.

Table 20. Space management for locally and remotely mounted file systems

Cluster A	Cluster B
A space management client is installed	A space management client is installed
File system A	File system A (from Cluster A remotely mounted on Cluster B)
Mounted locally	Mounted remotely
Space-managed by A	Not Space-managed by B

Only file systems of the local cluster can be managed by HSM. The local Cluster A file system is space-managed by Cluster A, but cannot be space-managed by Cluster B, although it is remotely mounted on Cluster B. The Data Management Application Programming Interface (DMAPI) interface does not work for remotely mounted file systems, so incorrect values might be reported by `dsmls` or `dsmdu`.

Syntax



Parameters

options

If you do not specify either of the following options, the space management client displays the number of 1 KB blocks that the specified directory and each of its subdirectories use.

-Allfiles

Displays the number of 1 KB blocks that each file in the specified directory and each of its subdirectories use.

-ERRORLOGName=*file_path*

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the `errorlogname` option in options file `dsm.sys` or `dsm.opt`, or as specified by the `DSM_LOG` environment variable. If any part of the path you specify does not exist, the space management client creates it.

-Help

Displays syntax and options for the command. Do not specify other options when you specify the `help` option.

-Summary

Displays only the total of 1 KB blocks that the specified directory and its subdirectories use.

directoryspec

The directory for which you want to display information. The default is the

current directory and its subdirectories. Use wildcard characters to specify more than one directory. You can issue more than one directory specification in one command. If you issue several directory specifications, separate each name with one or more blank spaces.

Examples

Task	Command
Display space usage information for the current directory and all of its subdirectories.	<code>dsmdu</code>
Display space usage information for the <code>/migfs3/test</code> directory and all of its subdirectories.	<code>dsmdu /migfs3/test</code>
Display space usage information for each file in the <code>/migfs2/test</code> directory and in all of its subdirectories.	<code>dsmdu -a /migfs2/test</code>
Display the total number of 1 KB blocks that the <code>/migfs2/test</code> directory and all of its subdirectories use.	<code>dsmdu -Summary /migfs2/test</code>

dsmls

The **dsmls** command displays file information, such as sizes and state.

Specifically, the **dsmls** command displays the following information about a list of files:

- Actual size (in bytes)
- Resident size (in bytes)
- Resident block size (in KB)
- File state and recall mode
- Name of the coupled IBM Spectrum Protect server
- File name

For a resident or premigrated file, the actual size and resident size are the same. For a migrated file, the actual size is the size of the original file. The resident size is the size of the stub file that remains on your local file system.

The file state for a file can have any of the following values: migrated (m), premigrated (p), or resident (r). A dash - indicates a directory or a non-regular file; for example, a character special file or a named pipe file. For a migrated file, the **dsmls** command also indicates the recall mode that you set for the file:

- If you set the recall mode to normal, additional information does not appear in the file state column.
- If you set the recall mode to partial file recall, the notation (p) displays in the file state column.
- If you set the recall mode to streaming, the notation (s) displays in the file state column.

Restriction:

- Running **dsmls** or **dsmdu** on remotely mounted GPFS file systems might show incorrect values. Consider the following example:

Table 21. Space management for locally and remotely mounted file systems

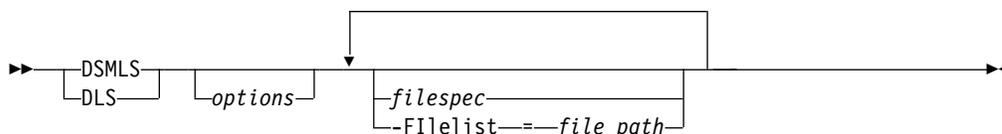
Cluster A	Cluster B
A space management client is installed	A space management client is installed

Table 21. Space management for locally and remotely mounted file systems (continued)

Cluster A	Cluster B
File system A	File system A (from Cluster A remotely mounted on Cluster B)
Mounted locally	Mounted remotely
Space-managed by A	Not Space-managed by B

Only file systems of the local cluster can be managed by HSM. The local Cluster A file system is space-managed by Cluster A, but cannot be space-managed by Cluster B, although it is remotely mounted on Cluster B. The Data Management Application Programming Interface (DMAPI) interface does not work for remotely mounted file systems, so incorrect values might be reported by **dsmls** or **dsmdu**.

Syntax



Parameters

options

-ERRORLOGName=*file_path*

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file `dsm.sys` or `dsm.opt`, or as specified by the **DSM_LOG** environment variable. If any part of the path you specify does not exist, the space management client creates it.

-Help

Displays syntax and options for the command. Do not specify other options when you specify the **help** option.

-Noheader

Omits column headings from the output for this command.

-Recursive

Displays information about files in subdirectories of the directory.

filespec

The path name for the files that you want to list. The default is all files in the current directory. Use wildcard characters to specify a group of files or all the files in a directory. You can enter more than one file specification in a command. If you enter several file specifications, separate each specification with one or more blank spaces.

-Filelist=*file_path*

Specifies the path of a file that contains a list of files to be processed by this command.

The entries in the list file must adhere to the following rules:

- Each entry is a fully qualified path to a file or a relative path to a file.

- No entry is a directory object.
- Each entry is on a separate line.

The space management client ignores any entry that does not adhere to these rules.

The following rules apply to a file list for the space management client. These rules differ from the rules for backup-archive client file lists.

- An entry can contain unprintable characters, but cannot contain a carriage return.
- Wildcard characters are allowed.
- A file path that contains blank spaces must be enclosed in quotation marks.
- An entry that begins with a quotation mark and ends with a quotation mark is tolerated. The space management client assumes that the quotation marks are not needed and omits them when the space management client processes the entry. If beginning and ending quotation marks are required to identify the file, add double quotation marks before the beginning and after the end of the entry.

The following example shows a list of files within a list file:

```
/home/dir/file1
"/fs1/dir2/file3"
"/fs2/my files/file4"
../tivoli/'file1'
"'fs3'/dir3/'file.txt'"
fs4/dir/a"file".txt
'/fs4/dir/file.txt'
/fs5/dir/file*with?wildcards.txt
```

If the list file name that you specify with the **filelist** option does not exist, the command fails.

Examples

Task	Command
List all files in the current directory.	<code>dsmls</code>
List all files in the <code>/migfs2/test</code> directory.	<code>dsmls /migfs2/test</code>
List all files in the <code>/migfs2/test</code> directory and in its subdirectories.	<code>dsmls -Recursive /migfs2/test</code>
List all files whose names begin with tf in the <code>/migfs2/test</code> directory.	<code>dsmls /migfs2/test/tf*</code>
List all files in the list file named <code>/tmp/filelist</code> .	<code>dsmls -filel=/tmp/filelist</code>

dsmmigfs add and update

Use the **dsmmigfs** command with the **add** parameter to add space management to the file system. Use the **dsmmigfs** command with the **update** parameter to update space management settings for your file system.

You must have root user authority to use this command.

For information about configuring GPFS integration with the space management client, see Technote 7018848.

Before you run `dsmmigfs add filesystem` on GPFS file systems, ensure that the file system is mounted and enabled for Data Management Application Programming Interface (DMAPI) management. Issue the following commands:

For GPFS file systems:

```
/usr/lpp/mmfs/bin/mm1sfs DevicePath -z
```

If it is required, change the value as follows:

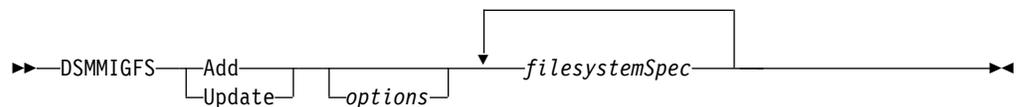
```
/usr/lpp/mmfs/bin/mmchfs DevicePath -z yes
```

Run only one **dsmmigfs** command within the local GPFS node group at the same time.

You cannot add space management to your root (/), /tmp, /usr, or /var file systems.

Note: You can query only information about mounted file systems. If a file system is space-managed, but not mounted, it does not show up within the query command

Syntax



Parameters

Add

Adds space management to your file systems.

Update

Updates one or more space management settings for a file system to which you added space management.

If you change the high and low thresholds or the premigration percentage, the new values take effect immediately. If you change the stub file size, the new size is used for files that are migrated after the change. The size of existing stub files does not change.

If you change the **minmigfilesize** option value, the new value is used the next time a migration candidates list is built. Also, the new value is used only for files that are migrated after the value is changed.

If you change the **readstartsrecall** option value, the new value is used only for files that are migrated after the value is changed.

You can set the quota to less than the amount of currently migrated and premigrated data. No more files are migrated until automatic and selective recalls reduce the amount of currently migrated and premigrated data to less than the new quota.

filesystemSpec

Specifies a file system name. You can specify more than one file system name, and you can use wildcard characters within a file system name. If you specify more than one file system name, separate each name with one or more blank spaces.

options

Use the option settings that are provided for this command to add or update space management settings.

On GPFS file systems, you can use the GPFS policy engine to monitor space thresholds and search for migration candidates. If you configure the GPFS policy engine to drive automatic migration, the following options for the **dsmmigfs** command have no consequence:

- **hthreshold**
- **lthreshold**
- **maxcandidates**
- **maxfiles**
- **minpartialrecallsize**
- **minmigfilesize**

-HThreshold=*n*

Specifies the high threshold percentage that you set for space usage on your file systems. Specify a value of 0 -100 percent. The default is 90 percent.

-Lthreshold=*n*

Specifies the low threshold percentage that you set for space usage on your file systems. Specify a value of 0 through 100 percent. The default is 80 percent.

-ERRORLOGName=*file_path*

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file **dsm.sys** or **dsm.opt**, or as specified by the **DSM_LOG** environment variable. If any part of the path you specify does not exist, the space management client creates it.

-INlinecopymode=*MIG|PREMIG|OFF*

Specifies how a backup operation handles a migrated file. This option applies only if backup operations are processed on migrated files.

If the **inlinecopymode** option is set to **MIG** and a migrated file is a candidate for a backup the file will be cloned inside the space management storage pool in the IBM Spectrum Protect server. The clone of the file will be stored in the backup storage pool in the IBM Spectrum Protect server. This operation is called inline copy.

If the **inlinecopymode** option is set to **PREMIG** and a premigrated or a migrated file is a candidate for backup the file the inline copy will happen.

If the **inlinecopymode** option is set to **OFF** and a migrated file is a candidate for backup the file will be recalled and the file state will change to premigrated. Then the backup processing will happen.

The inline copy function will be automatically disabled if one of the following takes place:

- The space management storage pool and the backup storage pool are not stored in the same IBM Spectrum Protect server instance.
- The file to be backed up has ACL or extended attribute metadata.

Note:

- The inline copy prevents data transfer between the IBM Spectrum Protect backup-archive client and the IBM Spectrum Protect server for backup processing and can be used if the file has a valid copy in the space management storage pool.
- IBM Spectrum Protect container storage pools do not support cloning. If container storage pools are used for space management data the `inlinecopymode` option must be set to OFF.

The default value is MIG.

-LOGname=*file_path*

Specifies the path and file name of the log file to be used by this command. The kinds of events that are logged to the file are specified with the `hsmlogeventflags` option. Error events are not logged to the file.

-Maxcandidates=*n*

Specifies the maximum number of migration candidates the `dsmscoutd` daemon delivers to the automatic migration process during automigration. A value 50 - 1000 is sufficient for good performance. The acceptable range is 9 - 9999999. The default is 10000.

-MAXFiles=*n*

Specifies the maximum number of files for which the CFI database is sized. The acceptable range is 0 - 4294967295. The default is 0.

If the specified value is less than the number of currently used blocks, an error occurs. The value must be greater than the current number of files in the file system. The value must be less than the total likely number of inodes.

You can configure the size of the complete file index (CFI). This value can be queried or set with the `maxfiles` parameter. If this parameter is set to 0, the CFI allocates the maximum required space. The maximum required space is the maximum number of blocks that are possible in the file system. Less file system space is needed to create the CFI. If the specified value is out of bounds, the nearest boundary is used to allocate the configurable CFI size.

-MINMigfilesize=*n*

Specifies the minimum size (in bytes) for a file to qualify for migration. The size is the space that a file occupies on the file system. Valid values are 0 - 2147483647. The default is 0. If you specify 0, the space management client uses the current file system stub size plus 1 byte as the minimum size for files that can be migrated. If you specify a nonzero value, the value must be greater than the current file system stub size. A valid value for a file system takes precedence over the global `minmigfilesize` option setting that is specified in the `dsm.sys` file.

-MINPartialrecallsize=*n*

Specifies the minimum size (in megabytes) for a file to qualify for partial file recall. The acceptable range is 0 - 999999999. The default value of 0 disables partial file recall for all files.

Valid only on GPFS file systems.

-MINStreamfilesize=*n*

Specifies the minimum file size for streaming recall mode. Streaming recall mode is set for files that are larger than the value of the `minstreamfilesize` option. Smaller files are recalled in normal recall mode.

Valid values are 0 - 999999999. The default is 0.

During a streaming recall, the recalled part of the file can be accessed before the complete file is recalled.

-Pmpercentage=*n*

Specifies the percentage of file system space that is available to contain premigrated files. The minimum value is 0. The maximum value is the low threshold. A warning is displayed if the premigration percentage is greater than the difference between the low threshold and the minimum size of the file system. A premigration percentage greater than this difference can result in insufficient space for the file system, and can result in endless attempts to find premigration candidates.

-PREViewsize=*n*

Specifies the preview size of the stub file. Valid values are 0 - 1,073,741,824, and must not be greater than the stub size.

When the attributes are set and the preview size is defined, a recall operation begins only when an application reads the stub file beyond the preview section of the stub file. If an operation reads only the preview section of the stub file, the file is not recalled.

-Quota=*n*

Specifies the maximum number of megabytes of data that you can migrate and premigrate from your file system to IBM Spectrum Protect storage. Specify a value from 0 to 999999999999999. The default is the number of megabytes that are allocated for your file system. If you set the quota to 0 for your file system, files do not migrate to storage. If you set the quota to 999999999999999, the amount of data you can migrate and premigrate is unlimited.

-READEVEnttimeout=*n*

Specifies the maximum time of inactivity (in seconds) before a streaming recall process times out. The acceptable range is 0 - 999999999. A value of 999999999 means that a streaming recall process does not time out. The default value is 600. This option applies to streaming recall mode only.

-REAdstartsrecall=No|Yes

Specifies whether a recall operation starts when an application reads the stub file.

This option applies only if streaming recall mode is set.

If the **previewsize** option is set for the stub file, a recall operation starts only when an application reads beyond the preview section of the stub file.

The default value is No.

Regardless of the value of the **readstartsrecall** option, the file is recalled when either of the following conditions are true:

- The recall daemon determines that the stub file does not contain all of the data that is required by a read operation.
- An operation modifies the file.

-SErver=*server_name*

Overrides the default migration server for this file system. Specify the server to contact for space management services. Define the server in a stanza in your `dsm.sys` file. If you do not specify a server name, the default migration server that you defined is used. Use a dash (-) to set the server to the default migration server.

-STREAMSeq=n

Specifies the number of megabytes that are buffered before the recall daemon flushes the data to disk. The acceptable range is 0 - 1024. The default value of 0 means that buffering is disabled. This option applies only to streaming recall mode.

-STUBSize=n

Specifies the size of stub files that remain on the file system when files are migrated to storage.

For the space management client on GPFS file systems, you can specify 0 or a multiple of the file system block size. The default value is 0.

For all file system types, the maximum value for a stub file size is 1 GB.

If HSM is globally deactivated on a node (**dsmmigfs globaldeactivate**), it is reset to active state, if one of the following commands are performed:

```
dsmmigfs add
dsmmigfs remove
dsmmigfs update
dsmmigfs takeover
dsmmigfs rollback
dsmmigfs globalreactivate
```

Examples

Commands are provided for the example tasks.

Task	Command
Add space management to the /hsmmanagedfs1 file system. Set the space to the maximum number of blocks in the file system.	dsmmigfs Add -MAXFiles=0 /hsmmanagedfs1
Update the space management settings for the /hsmmanagedfs2 file system. Set the space (number of blocks) to the specified number.	dsmmigfs update -MAXFiles=10000000 /hsmmanagedfs2
Add space management to the /home file system. Set the high threshold to 80 percent. Set the low threshold to 70 percent. Set the size of stub files to 256K (KB).	dsmmigfs Add -HT=80 -L=70 -STUBS=256k /home
Add space management to more than one file system and accept the default values for all space management settings.	dsmmigfs Add /home /test1 /proj*
Update the space management settings for the /home file system as follows: <ul style="list-style-type: none"> • Change the high threshold to 80 percent. • Change the low threshold to 70 percent. • Set the size of stub files to 1 megabyte. 	dsmmigfs Update -HT=80 -L=70 -STUBS=1m /home
Specify the minimum size of files in the /home/user1 file system that are recalled with partial file recall.	dsmmigfs Update -minp=100 /home/user1
Specify that the minimum size for a file that can be migrated from the /home/user2 file system is 1 megabyte.	dsmmigfs Update -minm=1048576 /home/user2

Task	Command
Add space management to the /myfs file system. Set the high threshold to 80 percent. Set the low threshold to 50 percent. Set the space (number of blocks) to the specified number.	<code>dsmmigfs add -maxfiles=1000000 -lt=50 -ht=80 -pm=10 /myfs</code>
When a file is migrated from the /myfs file system, create a stub file with the following specifications: <ul style="list-style-type: none"> The stub size is 32,768 bytes. The stub preview size is 16,384 bytes. If an application reads more than the preview size of the stub file, and the file is greater than 2 MB, recall the file. 	<code>dsmmigfs update -stubs=32768 -preview=16384 -minstreamfilesize=2 readstartsrecall=yes /myfs</code>

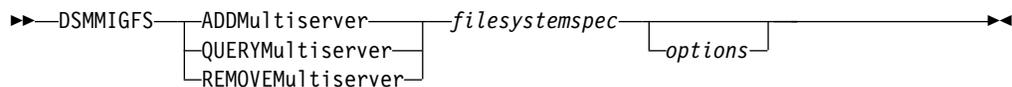
Related reference:

“`minmigfilesize`” on page 107

dsmmigfs addmultiserver, querymultiserver, and removemultiserver

Use the **dsmmigfs** command with the **addmultiserver**, **querymultiserver**, or **removemultiserver** parameter to manage space in an environment with multiple IBM Spectrum Protect servers.

Syntax



Parameters

ADDMultiserver

Adds one IBM Spectrum Protect server to the list of servers that can manage the file system.

QUERYMultiserver

Queries the status of the multiple-server environment. The command prints all IBM Spectrum Protect servers that can manage the specified file system. For each server, the output displays the number of files, number of bytes, and migration throughput.

REMOVEMultiserver

Removes an existing IBM Spectrum Protect server from the list of servers that can manage the specified file system. That IBM Spectrum Protect server is no longer a target for migration and backup. The current migration and backup processes are not affected. Removal does not affect the recall or restore of files that are migrated or backed up to the IBM Spectrum Protect server.

After you remove the IBM Spectrum Protect server from the list of servers, and before physically removing the server, run the `dsmRemoveServer.pl` script. The `dsmRemoveServer.pl` script recalls all files from the IBM Spectrum Protect server to the local file systems. After you run the `dsmRemoveServer.pl` script, you can migrate and back up files to another IBM Spectrum Protect server.

filesystemspec

The file system name. Specify only one file system.

options

-ERRORLOGName=*file_path*

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file `dsm.sys` or `dsm.opt`, or as specified by the **DSM_LOG** environment variable. If any part of the path you specify does not exist, the space management client creates it.

-Logname=*file_path*

Specifies the path and file name of the log file to be used by this command. The kinds of events that are logged to the file are specified with the **hsmlogeventflags** option. Error events are not logged to the file.

-Server=*server_name*

Specify the target server for the task. This option is valid only in a multiple-server environment. If the file system is managed by multiple servers and you do not specify **server**, the task is attempted with the default migration server.

If a file is coupled with a server, the value of this option must be that coupled server. If you specify another IBM Spectrum Protect server, the task fails.

Related tasks:

“Enabling a file system to be managed by multiple IBM Spectrum Protect servers” on page 46

“Removing IBM Spectrum Protect servers from a multiple-server environment” on page 48

Related reference:

“`dsmMultiServerUpgrade.pl`” on page 169

“`dsmRemoveServer.pl`” on page 171

dsmmigfs control

Use the **dsmmigfs** command with the **control** parameter to dynamically change space management client options applied from running space management daemons.

You must have root user authority to use this command.

Syntax

►► `DSMMIGFS Control` options ►►

Parameters

Control

Send a control statement to a space management client daemon to change the behavior of the daemon.

options

Use the option settings that are provided for this command to change client option settings applied from space management client daemons.

The control command applies for the space management client daemons watch daemon (dsmwatchd) and recall daemon (dsmrecalld).

-Type=enabletrace|disabletrace

The argument **Type** is mandatory. Specify whether the space management client tracing function should be enabled or disabled for a specified space management client daemon.

-PROcess=WD|RMD|RDD

The parameter **PROcess** is mandatory. Specify the space management client daemon process you want to update with new space management client option settings.

Specify the option **WD** to update the space management watch daemon (dsmwatchd).

Specify the option **RMD** to update the space management master recall daemon (dsmrecalld with PPID=1).

Specify the option **RDD** to update the space management distributor recall daemon (dsmrecalld with PPID=PID from master recall daemon).

-Options=name:value[:name:value]

The parameter **Options** is mandatory if the parameter **-Type** is set to enabletrace.

Specify the option **TRACEFILE** in combination with option value trace file name. Specify the option **TRACEFLAGS** in combination with the option value trace flags.

The trace file name must specify a file name in a directory that can be accessed from the space management client daemon process.

The remote node is not properly configured on the server.

Ensure that the remote node is properly configured and connected to the server by using TLS. This validates the remote node configuration and ensures that the information related to the remote node is sent to the server.

The trace flags must be one space management client trace flag or a list of space management client trace flags separated with blanks and enclosed with double quotes.

Note: In the current implementation of the control statement only trace options can be controlled for the specified space management client daemons.

Examples

Commands are provided for the example tasks.

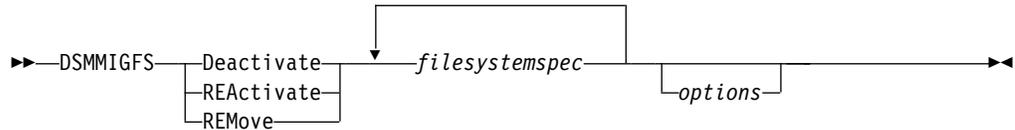
Task	Command
Enable tracing for the space management client daemon watch daemon with the /tmp/dsmwatchd.trc trace file and the SM SMVERBOSE SERVICE trace options.	dsmmigfs control -type=enabletrace -process=WD -options=tracefile:/tmp/dsmwatchd.trc::traceflags:"SM SMVERBOSE SERVICE"
Disable tracing for the space management client daemon recall master daemon.	dsmmigfs control -type=disabletrace -process=RMD

dsmmigfs deactivate, reactivate, and remove

Use the **dsmmigfs** command with the **deactivate**, **reactivate**, or **remove** parameter to deactivate, reactivate, or remove space management from a file system.

You must have root user authority to use this command.

Syntax



Parameters

Deactivate

Deactivates space management for a file system. The space management client cannot migrate or recall files, and cannot reconcile the file system. However, you can update space management settings for your file system, and access resident and premigrated files.

REActivate

Reactivates space management for a file system.

REMove

Removes space management from a file system. If you deactivated space management for your file system, reactivate it before you remove space management. If any orphaned stub files are located, the command fails. To remove space management, resolve all orphaned stub files, and issue the **dsmmigfs** command again.

fileSystemsSpec

The file system name. You can specify more than one file system name, and you can use wildcard characters within a file system name. If you specify more than one file system name, separate each name with one or more blank spaces.

options

-ERRORLOGName=*file_path*

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file `dsm.sys` or `dsm.opt`, or as specified by the **DSM_LOG** environment variable. If any part of the path you specify does not exist, the space management client creates it.

-Logname=*file_path*

Specifies the path and file name of the log file to be used by this command. The kinds of events that are logged to the file are specified with the **hsmlogeventflags** option. Error events are not logged to the file.

Examples

Task	Command
Deactivate space management for the /home file system.	<code>dsmmigfs Deactivate /home</code>

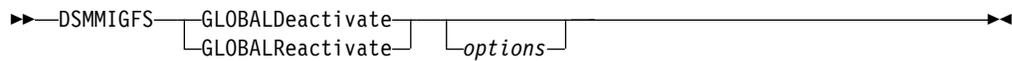
Task	Command
Reactivate space management for the /home file system.	<code>dsmmigfs REActivate /home</code>
Remove space management from the /home file system.	<code>dsmmigfs REMove /home</code>

dsmmigfs globaldeactivate and globalreactivate

Use the **dsmmigfs** command with the **globaldeactivate** or **globalreactivate** parameter to either deactivate or reactivate space management for a space-managed client node.

You must have root user authority to use this command.

Syntax



Parameters

GLOBALDeactivate

Deactivates space management for all file systems on your client node. The space management client cannot migration or recall any files and cannot reconcile any file system. However, you can update space management settings for file systems, add space management to other file systems, or access resident and premigrated files.

GLOBALReactivate

Reactivates space management for your client node. All file systems to which you added space management return to their previous state, including that which you added while space management was globally deactivated.

options

-ERRORLOGName=*file_path*

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file `dsm.sys` or `dsm.opt`, or as specified by the **DSM_LOG** environment variable. If any part of the path you specify does not exist, the space management client creates it.

-Logname=*file_path*

Specifies the path and file name of the log file to be used by this command. The kinds of events that are logged to the file are specified with the **hsmlogeventflags** option. Error events are not logged to the file.

Examples

Task	Command
Globally deactivate space management for your client node.	<code>dsmmigfs GLOBALDeactivate</code>
Globally reactivate space management for your client node.	<code>dsmmigfs GLOBALReactivate</code>

dsmmigfs help

Use the **dsmmigfs** command with the **help** parameter to display the **dsmmigfs** command syntax and options.

Syntax

►►—DSMMIGFS—Help—options—►►

Parameters

Help

Displays the syntax of the **dsmmigfs** command including command parameters, options, and valid ranges of options values.

options

-ERRORLOGName=*file_path*

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file `dsm.sys` or `dsm.opt`, or as specified by the **DSM_LOG** environment variable. If any part of the path you specify does not exist, the space management client creates it.

-Logname=*file_path*

Specifies the path and file name of the log file to be used by this command. The kinds of events that are logged to the file are specified with the **hsmlogeventflags** option. Error events are not logged to the file.

Examples

Task	Command
Display the syntax of the dsmmigfs command.	<code>dsmmigfs h</code>

dsmmigfs query

Use the **dsmmigfs** command with the **query** parameter to display the current space management settings for a file system.

Syntax

►►—DSMMIGFS—Query—options—↓
filesystemspec—►►

Parameters

Query

Display the current space management settings for the named file system.

options

-Detail

Display detailed HSM settings for the file system. Some settings are displayed only when queried with the **detail** option:

- The maximum number of migration candidates that are identified by the scout daemon
- The minimum partial file recall size
- The minimum streaming recall size
- The minimum size (in bytes) that qualifies a file for migration
- The local status for the distributed recall environment
- The status for the failover environment

For GPFS file systems only:

The **dsmmigfs query** command without the **detail** option displays only locally managed file systems. Use the **detail** option to display information for all space-managed file systems within the GPFS cluster.

When you specify the **detail** option, the HSM client also displays the node name for each node ID.

-Failover

Display a status overview of the failover environment of all HSM managed cluster nodes. Valid only on GPFS file systems. The output displays the status for the node name, and node ID. The following status can be displayed:

ENABLED

The node can take over space-managed file systems when another HSM node fails. If the node fails, control of the space-managed file systems fails over to other HSM nodes.

ENABLED_TAKEOVER

The node can take over space-managed file systems when another HSM node fails. If the node fails, control of the space-managed file systems does not fail over to other HSM nodes.

ENABLED_HANDOVER

If the node fails, control of the space-managed file systems fails over to other HSM nodes. The node does not take over space-managed file systems when another HSM node fails.

DISABLED

The HSM node does not take control of another HSM node's space-managed file systems. If the node fails, control of the space-managed file systems does not fail over to other HSM nodes.

To enable failover of HSM management of GPFS file systems on source nodes within a cluster environment, issue the **dsmmigfs enablefailover** command on each source node.

-ERRORLOGName=*file_path*

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file **dsm.sys** or **dsm.opt**, or as specified by the **DSM_LOG** environment variable. If any part of the path you specify does not exist, the space management client creates it.

-Logname=*file_path*

Specifies the path and file name of the log file to be used by this command. The kinds of events that are logged to the file are specified with the **hsmlogeventflags** option. Error events are not logged to the file.

-Node=*node_value* -parsable

Display the GPFS status of cluster nodes that are managed by the HSM client. Valid only on GPFS file systems. The minimum abbreviation of the option is one character (**-n=*node_value***). The output displays the following information for each node:

- GPFS node name
- GPFS node ID
- GPFS status. The GPFS status can be arbitrating, active, down, or unknown. The status is displayed when you specify the **detail** option.
- Recall daemon session ID
- Mount disposition
- Ping recall daemon
- Watch daemon session ID
- HSM status. The HSM status can be active or down.

Specify *node_value* by using one of three formats:

Specify *node_value* as a node name

Example: **-node=system1**

Specify *node_value* as a file that contains a list of node names

Example: **-node=/usr/tivoli/tsm/data/nodes.list**

Specify *node_value* by using a key word that defines a class of nodes

Specify one of the following node classes:

all All nodes in the GPFS cluster.

clientnodes

All nodes that do not participate in file system administration activities.

managernodes

All nodes in the pool of nodes from which file system managers and token managers are selected.

nonquorumnodes

All non-quorum nodes in the GPFS cluster.

nsdnodes

All NSD server nodes in the GPFS cluster.

quorumnodes

All quorum nodes in the GPFS cluster.

Example: **-node=quorumnodes**

When you specify the **node** option, you can specify the **parsable** option. The **parsable** option is optional. The **parsable** option displays output in a format that can be easily parsed. The minimum abbreviation of the option is four characters (**-pars**). The parsable format displays the following attributes in order:

- Command name
- Query type
- GPFS node name

GPFS node ID
 GPFS status
 HSM status
 Recall master session ID
 Mount disposition
 Recall daemon readiness
 Watch daemon session ID

Each attribute value ends with a colon.

The following output is an example of the format when you specify the **parsable** option:

```
dsmmigfs:queryNode:interceptor:2:active:active:51AC92F500000000:YES:YES:
519EF96D00000000:
```

filesystemspec

The file system name that displays current space management settings. The default is all space managed file systems.

Examples

Task	Command
Display the space management settings for the /migfs2 file system.	<code>dsmmigfs query /migfs2</code>
Display the space management entries that are made in the log file /tmp/migfslog.	<code>dsmmigfs query -L=/tmp/migfslog /migfs2</code>

Related concepts:

Chapter 3, “Configuring the space management client,” on page 19

Related reference:

“**dsmmigfs enablefailover**, and **disablefailover**” on page 143

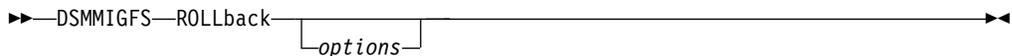
dsmmigfs rollback

Use the **dsmmigfs** command with the **rollback** parameter to transfer the space management of file systems to the preferred node if the node is different from the current owner node.

This command is valid only on GPFS file systems.

You must have root user authority to use this command.

Syntax



Parameters

ROLLback

Transfers the HSM management of file systems to the preferred node if the node is different from the current owner node. Enter this command on the preferred node.

options

-ERRORLOGName=*file_path*

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file `dsm.sys` or `dsm.opt`, or as specified by the **DSM_LOG** environment variable. If any part of the path you specify does not exist, the space management client creates it.

-Logname=*file_path*

Specifies the path and file name of the log file to be used by this command. The kinds of events that are logged to the file are specified with the **hsmlogeventflags** option. Error events are not logged to the file.

Examples

Task	Command
Transfer to the preferred node.	<code>dsmmigfs rollback</code>

dsmmigfs enablefailover, and disablefailover

Use the **dsmmigfs** command with the **enablefailover** or **disablefailover** parameter to manage recovery from partial system failure.

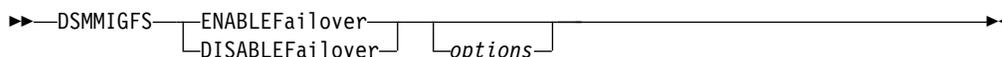
This command is valid only on GPFS file systems.

You must have root user authority to use this command.

One space management client can take over from an HSM client that is involved in a partial system failure if the following conditions are true:

- The failing space management client node has failover enabled.
- There are one or more additional space management client nodes within the same GPFS cluster with failover enabled.
- The space-managed file system is mounted on at least one of these nodes.
- A synchronous time exists on the failing nodes and the client nodes.
- The peer node is online

Syntax



Parameters

ENABLEFailover

Activates the node for failover operations within the GPFS cluster.

DISABLEFailover

Deactivates failover operations on the node.

options

-MODE=*value*

Specifies a failover mode. The **mode** option is valid only with the **enablefailover** parameter. You can specify the following values for the **mode** option:

TAKEOVER

The node can take over space-managed file systems when another HSM node fails. If the node fails, control of the space-managed file systems does not fail over to other HSM nodes.

HANDOVER

If the node fails, control of the space-managed file systems fails over to other HSM nodes. The node does not take over space-managed file systems when another HSM node fails.

ENABLED

The node can take over space-managed file systems when another HSM node fails. If the node fails, control of the space-managed file systems fails over to other HSM nodes.

The default is ENABLED

-ERRORLOGName=*file_path*

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file `dsm.sys` or `dsm.opt`, or as specified by the **DSM_LOG** environment variable. If any part of the path you specify does not exist, the space management client creates it.

-Logname=*file_path*

Specifies the path and file name of the log file to be used by this command. The kinds of events that are logged to the file are specified with the **hsmlogeventflags** option. Error events are not logged to the file.

Examples

Commands are provided for the example tasks.

Task	Command
Enable failover operations on the space management client node. If the node fails, control of the space-managed file systems fails over to other HSM nodes. The node does not take over space-managed file systems when another HSM node fails.	<code>dsmmigfs enablef -mode=HANDOVER</code>
Deactivate failover operations on the space management client node.	<code>dsmmigfs disableFailover</code>

dsmmigfs stop, start, and restart

use the **dsmmigfs** command with the **stop**, **start**, or **restart** parameter to control space management daemons.

Valid only on GPFS file systems.

You must have root user authority to use this command.

Use the **dsmmigfs** command with the **stop**, **start**, or **restart** parameter to control daemons:

- Start all daemons
- Stop all daemons, **dsmrecall** and **dsmmigrate** processes except **dsmwatchd**

- Restart all daemons, stop **dsmrecall** and **dsmmigrate** except **dsmwatchd**.

Note: Be aware that the daemons will be started with the same environment as the **dsmwatchd** daemon, which means that options files **dsm.opt** and **dsm.sys** in the default installation path **/usr/tivoli/tsm/client/ba/bin** will be used.

Syntax



Parameters

START

Starts all HSM daemons on the local client node. The **dsmwatchd** daemon is not affected.

STOP

Stops all HSM daemons. The **dsmrecall** and **dsmmigrate** processes are stopped. The **dsmwatchd** daemon is not affected.

RESTART

Restarts all HSM daemons. The **dsmrecall** and **dsmmigrate** processes are started. The **dsmwatchd** daemon is not affected.

options

-ERRORLOGName=*file_path*

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file **dsm.sys** or **dsm.opt**, or as specified by the **DSM_LOG** environment variable. If any part of the path you specify does not exist, the space management client creates it.

-Logname=*file_path*

Specifies the path and file name of the log file to be used by this command. The kinds of events that are logged to the file are specified with the **hsmlogeventflags** option. Error events are not logged to the file.

Examples

Task	Command
Start all daemons	dsmmigfs START
Stop all daemons	dsmmigfs STOP
Restart all daemons. For example, let them update the configuration set in your dsm.opt and dsm.sys options files.	dsmmigfs RESTART

dsmmigfs takeover

Use the **dsmmigfs** command with the **takeover** parameter to transfer the HSM management of a file system to another space management client node within the same local GPFS node set.

Valid only on GPFS file systems.

You must have root user authority to use this command.

Syntax

```
▶▶—DSMMIGFS—TAKEover—filespec—options—▶▶
```

Parameters

TAKEover

The **dsmmigfs** command transfers the HSM management of the specified file system to the space management client node on which you invoke this command. The transfer must be initiated on a node within the same local GPFS node set.

filespec

The name of the file system you want to takeover.

options

-ERRORLOGName=*file_path*

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file `dsm.sys` or `dsm.opt`, or as specified by the **DSM_LOG** environment variable. If any part of the path you specify does not exist, the space management client creates it.

-Logname=*file_path*

Specifies the path and file name of the log file to be used by this command. The kinds of events that are logged to the file are specified with the **hsmlogeventflags** option. Error events are not logged to the file.

Examples

Task	Command
Transfer the HSM management of the current directory to the space management client node within the same local GPFS node set.	<code>dsmmigfs takeover /home/filesystem</code>

-Detail

Use with the **mgmtclass** option to display information about each available management class. If you do not use this option, the HSM client displays the management class name and a brief description only.

-ERRORLOGName=*file_path*

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file *dsm.sys* or *dsm.opt*, or as specified by the **DSM_LOG** environment variable. If any part of the path you specify does not exist, the space management client creates it.

-Help

Displays syntax and options for the command. Do not specify other options when you specify the **help** option.

-Mgmtclass

Displays information about each management class that you can assign to your files.

-Options

Displays the current settings for your client and server options. This option is the default.

-Server=*server_name*

Specify the target server for the task. This option is valid only in a multiple-server environment. In a multiple-server environment you must specify the **server** option.

This option is valid with the **sortedall** and **sortedmigrated** options.

-SORTEDAll

Lists all files in the file system in this order: resident files, premigrated files, migrated files. Sorts migrated files in the most efficient order for recall.

-SORTEDMigrated

Lists all files that you migrated from the file system to IBM Spectrum Protect storage in the most efficient order for recall.

filesystemspec

The file system for which you want to display information. The default is the current file system. You can specify more than one file system name, and you can use wildcard characters within a file system name. If you specify more than one file system name, separate each name with one or more blank spaces.

Examples

Task	Command
Display the current settings for the client and server options.	<code>dsmmigquery</code>
Display information about management classes that you can assign to files on your client node.	<code>dsmmigquery -Mgmtclass -Detail</code>

HSM and backup-archive client `dsmmigquery` command shared options

The `dsmmigquery -o` command displays both HSM options and many options that are shared between the space management client and the backup-archive client.

The following is a list of the options that are shared between the space management client and the backup-archive client.

- `asnodename`
- `commethod`
- `compression`
- `defaultserver`
- `detail`
- `diskbuffsize`
- `enablelanfree`
- `errorlogmax`
- `errorlogname`
- `errorlogretention`
- `exclude`
- `exclude.compression`
- `includexcl`
- `include`
- `include.compression`
- `lanfreecommmethod`
- `lanfreetcpport`
- `lanfreeshmport`
- `makesparsefile`
- `nodename`
- `passwordaccess`
- `passworddir`
- `servername`
- `shmport`
- `skipacl`
- `tcpbuffsize`
- `tcpnodelay`
- `tcpport`
- `tcpserveraddress`
- `tcpwindowsize`

Related information:

 [Backup-archive client processing options](#)

`dsmmigrate`

The `dsmmigrate` command selects specific files from the local file system and migrates them to the IBM Spectrum Protect server.

Tip: On large file systems, selective migration can take a long time. Migration can be quicker if you premigrate files before you migrate them. The quickest way to create free space requires two steps:

1. Prepare for migration by premigrating files.
2. When you need to create free space quickly, use the `stubmigrated` option to migrate only premigrated files.

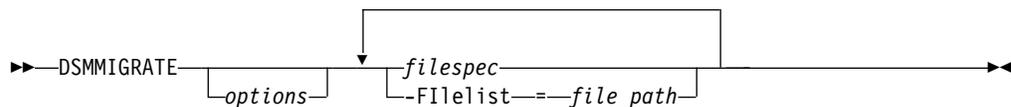
If the file exists in a file system with a different server stanza than the last processed file, a new session starts for each migrated file. This action can happen as a result of links from one file system to another.

Note:

1. You are prompted either to wait for the medium to mount or to skip the file when the following two conditions are met:
 - The **tapeprompt** option is set to YES in the dsm.opt file.
 - The destination for the migrated file is a storage pool that consists of removable media (such as tape)
2. The space management client does not migrate contents of symbolic links. Symbolic links are not followed during recursive selective migration.

The first file migrates even if the file size exceeds the quota that you specified for the file system. When migration occurs, the **ddf** command displays zero migrated and premigrated bytes for the file system. If the total number of bytes exceeds the quota after the file migrates, the next file is not migrated.

Syntax



Parameters

options

-Detail

Displays the size and file name for each file that you migrate.

-ERRORLOGName=*file_path*

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file dsm.sys or dsm.opt, or as specified by the **DSM_LOG** environment variable. If any part of the path you specify does not exist, the space management client creates it.

-Help

Displays syntax and options for the command. Do not specify other options when you specify the **help** option.

-Logname=*file_path*

Specifies the path and file name of the log file to be used by this command. The kinds of events that are logged to the file are specified with the **hsmlogeventflags** option. Error events are not logged to the file.

-Premigrate

A copy of the file is sent to IBM Spectrum Protect storage, and the file remains on the local file system. The file state changes to premigrated. You cannot use the **premigrate** option and the **stubpremigrated** option at the same time.

-PREViewsize=*n*

Specifies the preview size of the stub file. Valid values are 0 - 1,073,741,824, and must not be greater than the stub size.

When the attributes are set and the preview size is defined, a recall operation begins only when an application reads the stub file beyond the preview section of the stub file. If an operation reads only the preview section of the stub file, the file is not recalled.

-REAdstartsrecall=No|Yes

Specifies whether a recall operation starts when an application reads the stub file.

This option applies only if streaming recall mode is set.

If the **previewsize** option is set for the stub file, a recall operation starts only when an application reads beyond the preview section of the stub file.

The default value is No.

Regardless of the value of the **readstartsrecall** option, the file is recalled when either of the following conditions are true:

- The recall daemon determines that the stub file does not contain all of the data that is required by a read operation.
- An operation modifies the file.

-Recursive

Migrates files in any subdirectory of the specified directory that matches the file specification. If you do not use this option, only those files from the directories that you specify are migrated.

The space management client does not migrate contents of symbolic links. Symbolic links are not followed during recursive selective migration.

-Server=server_name

Specify the target server for the task. This option is valid only in a multiple-server environment. If the file system is managed by multiple servers and you do not specify **server**, the task is attempted with the default migration server.

If a file is coupled with any IBM Spectrum Protect server, you must specify that server. If a file is not coupled with any server, you can specify any IBM Spectrum Protect server that was added to the multiple-server environment.

-STUBSize=n

Specifies the size of stub files that remain on the file systems when files are migrated to storage.

For the space management client on GPFS file systems, you can specify 0 or a multiple of the file system block size. The default value is 0.

For all file system types, the maximum value for a stub file size is 1 GB.

-STUBPremigrated

Only files that are in premigrated state are migrated. The migration process yields free space on the file system quickly. File copies exist in IBM Spectrum Protect storage. The migration process only replaces the local files with stub files.

filespec

The path and file name of the files that you want to migrate. This parameter is required. You can use wildcard characters to specify a group of files or all files in a directory. You can enter more than one file specification in one command. If you enter more than one file specification, separate each specification with one or more blank spaces.

-filelist=*file_path*

Specifies the path of a file that contains a list of files to be processed by this command.

The entries in the list file must adhere to the following rules:

- Each entry is a fully qualified path to a file or a relative path to a file.
- No entry is a directory object.
- Each entry is on a separate line.

The space management client ignores any entry that does not adhere to these rules.

The following rules apply to a file list for the space management client. These rules differ from the rules for backup-archive client file lists.

- An entry can contain unprintable characters, but cannot contain a carriage return.
- Wildcard characters are allowed.
- A file path that contains blank spaces must be enclosed in quotation marks.
- An entry that begins with a quotation mark and ends with a quotation mark is tolerated. The space management client assumes that the quotation marks are not needed and omits them when the space management client processes the entry. If beginning and ending quotation marks are required to identify the file, add double quotation marks before the beginning and after the end of the entry.

The following example shows a list of files within a list file:

```
/home/dir/file1
"/fs1/dir2/file3"
"/fs2/my files/file4"
../tivoli/'file1'
"'fs3'/dir3/'file.txt'"
fs4/dir/a"file".txt
'/fs4/dir/file.txt'
/fs5/dir/file*with?wildcards.txt
```

If the list file name that you specify with the **filelist** option does not exist, the command fails.

Examples

Task	Command
Migrate all files in a directory that is named /migfs2/test/dir1 and in all of its subdirectories. Display the information.	<pre>dsmmigrate -Recursive -Detail /migfs2/test/dir1</pre>
Migrate a file that is named tf04 from the current directory and display the information. Log the results in the file /tmp/miglog.	<pre>dsmmigrate -Detail -L=/tmp/miglog tf04</pre>

Task	Command
Migrate all files in a file list that is named /tmp/filelist. Log the results in the file /tmp/miglog. Set the stub size to 4096 bytes and the preview size to 2048 bytes. Restriction: When a file is recalled, the preview size affects recall processing only if streaming recall mode is set and readstartsrecall=yes. You can set streaming recall mode with the minstreamfilesize option on the dsmmigfs command or with the recallmode option on the dsmattr command.	<code>dsmmigrate -logname=/tmp/miglog -filelist=/tmp/filelist -stubsizesize=4096 -previewsize=2048 -readstartsrecall=yes</code>
Use a shell application to build a list of files that are owned by user ibm, and then migrate the files.	<code>find /hsmfilesystem -user ibm -print > /tmp/filelist dsmmigrate -filelist=/tmp/ filelist</code>
Premigrate the file /migfs2/test/file1, but do not replace the local file with a stub.	<code>dsmmigrate -p /migfs2/test/file1</code>
Migrate all premigrated files in directory /migfs2/test/dir1/ and in all of the subdirectories	<code>dsmmigrate -R -ST /migfs2/test/dir1</code>

dsmmigundelete

The **dsmmigundelete** command recreates deleted stub files for migrated files, and creates stub files for premigrated files for which an original file does not exist on your local file system. The file then becomes a migrated file.

You must have root user authority to use this command.

When a stub file or an original copy of a premigrated file is deleted from your local file system, the corresponding migrated or premigrated file is marked for expiration when reconciliation runs again.

Note: The **dsmmigundelete** command does not support hard-linked files. If you recreate a stub file for a hard-linked file, a stub file is not recreated *unless* all of the hard-linked files are deleted from your local file system. When one file in a set of hard-linked files is migrated, all of the hard-linked files in the set become stub files. When the **dsmmigundelete** command recreates a stub file for a hard-linked file, the stub file has the same name as the file that was originally migrated. Stub files are not recreated for any other files that were previously in the hard-linked set of files.

Syntax

```

▶▶—DSMMIGUNDELETE—┬───┬───▶
                    │   │
                    └───┴───┬───┬───▶
                            │   │
                            └───┴───▶
                    [filesystemspec] [options]

```

Parameters

options

-Recover

Recreates stub files that have been removed from your file system. Use this option if you need better performance during the recover process.

Important: This option recreates all of the existing stub files that overwrites all premigrated or recalled files. Data is lost from the last migration up to when this option is used! After the **dsmmigundelete -recover** command is issued, there is no way you can check whether the data of the local stub file is different from the contents of the file in the IBM Spectrum Protect server.

Use this option only if all or most of the files in the file system are stub files and you need good performance during stub file recreation.

For performance purposes, do not use this option with the **detail** option.

Files that were never migrated are not recreated.

-Detail

Displays detailed information about which stub file is being recreated to the file system.

Do not use this option with the recover option.

-Expiring

Recreates a stub file for a migrated file if a corresponding stub file does not exist on your local file system, whether the migrated file was marked for expiration or not. Or, it creates a stub file for a premigrated file if a corresponding original file does not exist on your local file system, whether the premigrated file was marked for expiration or not.

Issue the **dsmmigundelete** command with the **expiring** option if you ran reconciliation since the files were deleted.

If you do not use the **expiring** option, the HSM client recreates a stub file for a migrated file if a corresponding stub file does not exist on your local file system and the migrated file was not marked for expiration. Or, it creates a stub file for a premigrated file if a corresponding original file does not exist on your local file system, and the premigrated file was not marked for expiration.

Enter the **dsmmigundelete** command *without* the **expiring** option if you did not run reconciliation since the files were deleted.

-ERRORLOGName=file_path

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file **dsm.sys** or **dsm.opt**, or as specified by the **DSM_LOG** environment variable. If any part of the path you specify does not exist, the space management client creates it.

-Logname=file_path

Specifies the path and file name of the log file to be used by this command. The kinds of events that are logged to the file are specified with the **hsmlogeventflags** option. Error events are not logged to the file.

-Server=server_name

Specify the target server for the task. This option is valid only in a multiple-server environment. If the file system is managed by multiple servers and you do not specify **server**, the task is attempted with the default migration server.

If a file is coupled with a server, the value of this option must be that coupled server. If you specify another IBM Spectrum Protect server, the task fails.

filesystemspec

The name of the file system for which you want to recreate deleted stub files and create stub files for premigrated files that were deleted from your local file system. The default is all file systems for which space management is active. You can specify more than one file system name. If you specify several file system names, separate each name with one or more blank spaces.

Examples

Task	Command
Recreate stub files that were accidentally deleted from the /home file system. Recreate stub files regardless of the expiration status of the migration copies (use the expiring option). Reconciliation was run since the files were deleted. Results are logged in the log file /tmp/undeletelog.	<code>dsmmigundelete -expiring -Logname=/tmp/undeletelog /home</code>
Recreate stub files for migrated files in the /home file system. Recreate stub files only for those migration copies that are not marked for expiration (do not use the expiring option). Reconciliation was not run because the files were deleted. Entries are made in the log file /tmp/undeletelog.	<code>dsmmigundelete -L=/tmp/undeletelog /home</code>
Recreates stub files for migrated files that had been removed from the /trullofs file system.	<code>dsmmigundelete -recover /trullofs</code>

dsmonitord

The **dsmonitord** command starts the HSM space monitor daemon if it has stopped. If you issue this command and the space monitor daemon is running, action is not taken.

You must have root user authority to use this command.

Syntax

```

▶▶—DSMMONITORD—┐
                   └options┘
  
```

Parameters

options

-ERRORLOGName=*file_path*

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file `dsm.sys` or `dsm.opt`, or as

Examples

Task	Command
Display the status of recall processes.	dsmq

dsmrecall

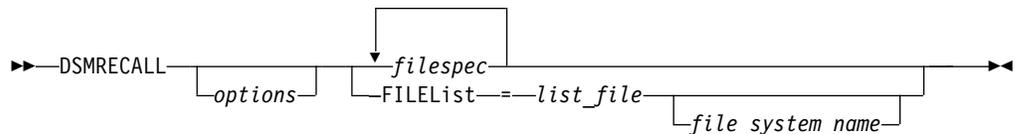
The **dsmrecall** command selectively recalls migrated files or parts of migrated files to the local file system. Space management must be active.

Note: On large file systems, selective recall can take a long time.

If the file resides in a file system whose server stanza is different from the stanza of the last processed file, a new session is started for each recalled file. This process can happen as a result of links from one file system to another.

To display a list of all migrated files, use the **dsmmigquery** command. To display information about a list of migrated files from a particular file system or directory, use the **dsmis** command.

Syntax



Parameters

options

-Detail

Display the size, path, and file name for each file that you recall.

-ERRORLOGName=*file_path*

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file `dsm.sys` or `dsm.opt`, or as specified by the **DSM_LOG** environment variable. If any part of the path you specify does not exist, the space management client creates it.

-Help

Displays syntax and options for the command. Do not specify other options when you specify the **help** option.

-OFFset=*n*

Specify the offset from the beginning of the file of the required data range for partial recall (in bytes, kilobytes, megabytes, or gigabytes). Multipliers (k, m, g, K, M, or G) can be used. The range of acceptable values is 0 - 2147483647. There is no default value.

Valid only on GPFS file systems.

Remember: Use this option only with the **size** option and when a file system is specified. The **recursive**, **detail**, and **filelist** options are not valid with the **offset** option.

Examples are: **-offset=10** (bytes), **-offset=23k** (kilobytes), **-off=5M** (megabytes), **-off=2G** (gigabytes).

-PREView

Generate list files that are optimized for tape recalls but do not recall the files. You must also specify **filelist** and a file system. The **preview** option is not valid when **filelist** specifies a collection file.

-Logname=*file_path*

Specifies the path and file name of the log file to be used by this command. The kinds of events that are logged to the file are specified with the **hsmlogeventflags** option. Error events are not logged to the file.

-Recursive

Recall files that match the file specification in a directory and its subdirectories. If you do not use this option, files are recalled only for those directories that you specify.

-RESident

When files are recalled, change the state to resident. If you do not use the **resident** option, the recalled files are in the premigrated state.

Remember: During the next reconciliation of the file system, files in resident state are marked for expiration on the IBM Spectrum Protect server.

-SIZE=n

Specify the size of the required data range for partial recall (in bytes, kilobytes, megabytes, or gigabytes). Multipliers (k, m, g, K, M, or G) can be used. The range of acceptable values is 0 - 4294967295. There is no default value.

Valid only on GPFS file systems.

Remember: Use this option only with the **offset** and the **filespec** options. The **recursive**, **detail**, and **filelist** options are not valid with the **size** option.

Examples are: **-size=10** (bytes), **-size=23k** (kilobytes), **-size=5M** (megabytes), **-size=2G** (gigabytes).

-Server=*server_name*

Specify the target server for the task. This option is valid only in a multiple-server environment. If the file system is managed by multiple servers and you do not specify **server**, the task is attempted with the default migration server.

If a file is coupled with a server, the value of **server** must be that coupled server. If you specify another IBM Spectrum Protect server, the task fails with a warning message.

filespec

Specify the path and file name of the files to recall. This parameter is required. You can use wildcard characters to specify a group of files or all files in a directory, or more than one file specification in one command. When you use wildcard characters in a file specification, the space management client recalls all files that match the specification. If a file matches the specification but it is

not migrated, an error message displays. If you enter more than one file specification, separate each specification with at least one blank space.

-FILEList=*list_file*

Process the files that are listed in *list_file*.

The entries in the list file must adhere to the following rules:

- Each entry is a fully qualified path to a file or a relative path to a file.
- No entry is a directory object.
- Each entry is on a separate line.

The space management client ignores any entry that does not adhere to these rules.

The following example shows a list of files within a list file:

```
/home/dir/file1
"/fs1/dir2/file3"
"/fs2/my files/file4"
../tivoli/'file1'
"'fs3'/dir3/'file.txt'"
fs4/dir/a"file".txt
'/fs4/dir/file.txt'
/fs5/dir/file*with?wildcards.txt
```

If the list file name that you specify with the **filelist** option does not exist, the command fails.

You can specify a collection file generated by **dsmrecall** or a list file generated by another application. You can specify only one list file or collection file, and you can specify the **filelist** option only once in each command.

If you also specify the preview option:

You must also specify *file_system_name*. The space management client generates tape-optimized list files and a collection file. No files are recalled.

If you do not specify the preview option:

The space management client recalls files.

If you also specify *file_system_name*:

The space management client performs a tape-optimized recall.

If the list file is a collection file that was generated by **dsmrecall:**

Recall starts. The list files referenced by the collection are processed in the order defined in the collection file. You must specify the collection file generated by **dsmrecall**; you cannot specify a tape list file or disk list file.

If the list file was not generated by **dsmrecall:**

The space management client first reorders the list to optimize recalls from tape. Then the space management client recalls the files.

If you do not specify *file_system_name*:

The space management client recalls the files in the list file. The space management client does not reorder the list to optimize recalls from tape.

file_system_name

Specify a file system. You must also specify the **filelist** option. The space management client optimizes the processing of files stored on tape.

Examples

Task	Command
Recall a single file named /migfs1/test/tf04 and display detailed information.	<code>dsmrecall -Detail /migfs1/test/tf04</code>
Recall all migrated files in a directory named /mfs4/user1 and all migrated files in its subdirectories.	<code>dsmrecall -Recursive /mfs4/user1/*</code>
Recall all files in the FILElist named /tmp/filelist. Results are logged in the log file /tmp/recalllog.	<code>dsmrecall -Logname=/tmp/recalllog -filelist=/tmp/filelist</code>
Recall the specified portion of a file named /mfs1/file10 in partial recall mode.	<code>dsmrecall -offset=10M -size=500M /mfs1/file10</code>

Related concepts:

"List files for optimized tape recalls" on page 76

Related reference:

"**dsm1s**" on page 126

"**dsmmigquery**" on page 147

dsmrecalld

The **dsmrecalld** command starts a recall daemon if it is not running.

You must have root user authority to use this command.

Note:

- No action is taken if you issue this command while a recall daemon is running.
- When using the backup-archive client to restore space-managed files, **dsmrecalld** must be running.

Syntax

→ `DSMRECALLD` options →

Parameters

options

-ERRORLOGName=*file_path*

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file `dsm.sys` or `dsm.opt`, or as specified by the **DSM_LOG** environment variable. If any part of the path you specify does not exist, the space management client creates it.

-Help

Displays syntax and options for the command. Do not specify other options when you specify the **help** option.

dsmreconcile

The **dsmreconcile** command synchronizes the file systems on your client node with the IBM Spectrum Protect server that you contact for space management services. Use this command at any time to reconcile one or more file systems. Specify how often to automatically reconcile with the **reconcileinterval** option and how long to keep obsolete copies with the **migfileexpiration** option in the `dsm.sys` file.

You must have root user authority to use this command.

Note: The **dsmreconcile** command must be found with the `PATH` variable, or the `dsmmonitord` daemon cannot reconcile the file system.

Syntax

```
▶—DSMRECONCILE—┬── options ─┬──▶
```

Parameters

options

-Detail

Prints progress messages.

-ERRORLOGName=*file_path*

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file `dsm.sys` or `dsm.opt`, or as specified by the **DSM_LOG** environment variable. If any part of the path you specify does not exist, the space management client creates it.

-FILEINFO

Expires and deletes obsolete copies from the IBM Spectrum Protect server. This option is the basic reconciliation task and is the default option.

-FILELIST=*file_path*

Specifies the path of a file that contains a list of files to be processed by this command.

-Help

Displays syntax and options for the command. Do not specify other options when you specify the **help** option.

-ORPHANCHECK

Checks for orphaned files in the local file system.

The HSM client verifies that valid objects for each local stub file exist in the space management pool of the IBM Spectrum Protect server. The **dsmreconcile** process determines all migrated and premigrated files, and checks that corresponding objects exist on the server. When orphans are located, their names are recorded in the *file_system_name*/.SpaceMan/orphan.stubs file. This command option overrides the value of the **checkfororphans** option in the `dsm.sys` file.

If you previously ran the **dsmreconcile** command with the **preptwo** option, IBM Spectrum Protect reconciles with a two-way orphan check.

The two-way orphan check processing identifies orphans on the file system and on the IBM Spectrum Protect server. The two-way check is done in a single pass and the orphan identification process uses parallel processing.

Tip: If reconciliation is controlled by the space management client space monitor daemon, the scout daemon must be running when you specify this option. If reconciliation is controlled by the GPFS policy engine, the scout daemon is not necessary.

When the reconciliation process is running in orphan-check mode, metadata information is updated on the server. Obsolete copies of files are not expired or deleted on the server.

-PREPTWO

This option prepares a file system for a two-way orphan check. Use this option when reconciliation is controlled by the GPFS policy engine.

When reconciliation is controlled by the GPFS policy engine, migrated and premigrated files are associated with DMAPI attribute **extObjId**. If you set **hsmextobjidattr=YES**, the DMAPI attribute **extObjId** is associated with the files when they are migrated or premigrated. The **preptwo** option makes the association for any files that were not associated with the **extObjId** attribute when the files were migrated or premigrated.

The **preptwo** option is valid only if **hsmextobjidattr=YES**.

When you specify the **preptwo** option, do not specify any of the following options:

- **fileinfo**
- **filelist**
- **orphancheck**

-Server=server_name

Specify the target server for the task. This option is valid only in a multiple-server environment. If the file system is managed by multiple servers and you do not specify **server**, the task is attempted with the default migration server.

If a file is coupled with a server, the value of this option must be that coupled server. If you specify another IBM Spectrum Protect server, the task fails.

filesystemspec

The name of the file system to reconcile. If you do not specify a file system name, the space management client reconciles all file systems on your workstation for which space management is active. If you enter more than one file system name, separate each name with at least one blank space.

Examples

Task	Command
Start reconciliation for all file systems for which space management is active.	<code>dsmreconcile</code>
Start reconciliation for the <code>/migfs1</code> file system.	<code>dsmreconcile /migfs1</code>
Start an orphan check reconciliation process for the <code>/migfs1</code> file system.	<code>dsmreconcile -orphancheck /migfs1</code>
Start reconciliation for file systems <code>/home</code> and <code>/test1</code> .	<code>dsmreconcile /home /test1</code>

Task	Command
Prepare the two-way orphan check reconciliation for the /migfs1 file system.	<code>dsmreconcile -preptwo /migfs1</code>
Start the policy-based reconciliation on the files that are listed in the /home/dsn/filelist file list.	<code>dsmreconcile -orphancheck -filelist=/home/dsn/filelist</code>

Related concepts:

“The scout daemon” on page 84

dsmrm

The **dsmrm** command removes a recall process from the queue. To obtain the required recall ID to remove a recall process, use the **dsmq** command.

You must have root user authority to use this command.

Note: After a recall process starts, issue the **dsmrm** command to stop the process. Do not use **Ctrl+C** to stop a recall process.

Syntax



Parameters

recallid

The recall process ID that you want to remove from the queue.

options

-ERRORLOGName=*file_path*

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file `dsm.sys` or `dsm.opt`, or as specified by the **DSM_LOG** environment variable. If any part of the path you specify does not exist, the space management client creates it.

-Help

Displays syntax and options for the command. Do not specify other options when you specify the **help** option.

Examples

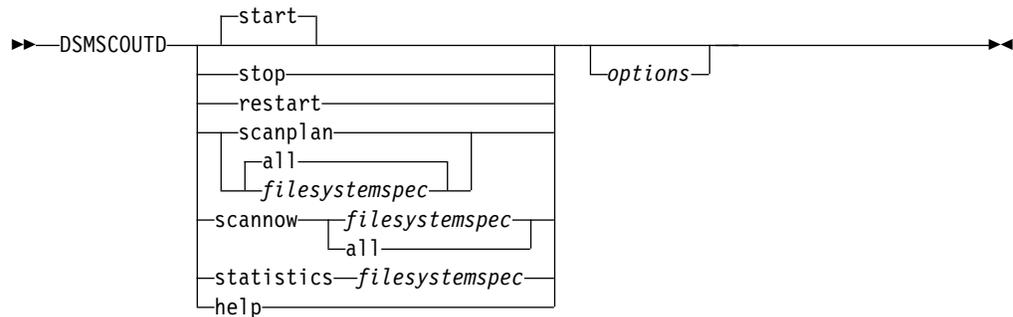
Task	Command
Remove recall ID 10 from the queue.	<code>dsmrm 10</code>
Remove recall ID 5 and recall ID 6 from the queue.	<code>dsmrm 5 6</code>

dsmscoutd

The **dsmscoutd** command starts, stops, and restarts the scout daemon. You can also use the command to start a scan, show future scan times, and show past scan information for file systems.

You must have root user authority to use this command.

Syntax



Parameters

All parameters must be entered as indicated. The parameters must be entered in lowercase. The parameters have no abbreviation. The parameters have no short format.

start

Starts the scout daemon. This option is the default, when no parameters are specified.

stop

Stops the scout daemon.

restart

Stops and restarts the daemon.

scanplan

Displays information about the next scan time, include the remaining time until the next scan, for one or more file systems. If the results of this command do not show the status of space-managed file systems, the dsmscoutd daemon lost its communication paths. Communications paths can be lost because of unexpected file system issues such as unmounting a file system. Use the **dsmscoutd restart** command to recover the communication paths for the daemon.

all

Use this option to include all file systems in the specific action.

filesystemspec

The file system name. You can specify more than one file system name, and you can use wildcard characters within a file system name. If you specify more than one file system name, separate each name with one or more blank spaces.

scannow

Starts scanning of one or more file systems.

statistics

Provides statistics about the current (if there is one) and last scan of the file system.

options

-ERRORLOGName=*file_path*

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file `dsm.sys` or `dsm.opt`, or as specified by the **DSM_LOG** environment variable. If any part of the path you specify does not exist, the space management client creates it.

-Help

Displays syntax and options for the command. Do not specify other options when you specify the **help** option.

dsmsetpw

The **dsmsetpw** command changes the IBM Spectrum Protect password for your client node.

You must have root user authority to use this command.

To keep your password secure, issue the **dsmsetpw** command without your old password and new password. The system prompts you for each one. When you are prompted to enter your old and new passwords, you eliminate the possibility that another user can display your password.

If your IBM Spectrum Protect server is at version 6.3.3 or later, and if you use an LDAP directory server to authenticate passwords

Use any of the following characters to create a password:

```
a b c d e f g h i j k l m n o p q r s t u v w x y z  
A B C D E F G H I J K L M N O P Q R S T U V W X Y Z  
0 1 2 3 4 5 6 7 8 9 . ! @ # $ % ^ & * _ - + = ` ( )  
| { } [ ] : ; < > , ? / ~
```

Passwords are case-sensitive and are subject to more restrictions that can be imposed by LDAP policies.

If your IBM Spectrum Protect server is at version 6.3.3 or later, and if you do not use an LDAP directory server to authenticate passwords

Use any of the following characters to create a password:

```
a b c d e f g h i j k l m n o p q r s t u v w x y z  
A B C D E F G H I J K L M N O P Q R S T U V W X Y Z  
0 1 2 3 4 5 6 7 8 9 . ! @ # $ % ^ & * _ - + = ` ( )  
| { } [ ] : ; < > , ? / ~
```

Passwords are stored in the IBM Spectrum Protect server database and are not case-sensitive.

If your IBM Spectrum Protect server is earlier than version 6.3.3

Use any of the following characters to create a password:

```
a b c d e f g h i j k l m n o p q r s t u v w x y z  
A B C D E F G H I J K L M N O P Q R S T U V W X Y Z  
0 1 2 3 4 5 6 7 8 9  
_ - & + .
```

Passwords are stored in the IBM Spectrum Protect server database and are not case-sensitive.

On the command line, enclose a password in single quotation marks (') if the password contains one or more special characters. The quotation marks ensure that the special characters are correctly interpreted as password characters. Without quotation marks, the special characters can be interpreted as shell escape characters, file redirection characters, or other characters that have significance to the operating system.

Syntax

```
▶▶—DSMSETPW—┌──────────┐ ┌──────────┐ ┌──────────┐ ┌──────────┐▶▶  
                options    —oldpassword— —newpassword—
```

Parameters

options

-ERRORLOGName=*file_path*

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file `dsm.sys` or `dsm.opt`, or as specified by the **DSM_LOG** environment variable. If any part of the path you specify does not exist, the space management client creates it.

-Help

Displays syntax and options for the command. Do not specify other options when you specify the **help** option.

oldpassword

The current IBM Spectrum Protect password for your client node.

newpassword

The new IBM Spectrum Protect password to set for your client node.

Examples

Task	Command
Change your current IBM Spectrum Protect password from <code>osecret</code> to <code>nsecret</code> .	<code>dsmsetpw osecret nsecret</code>

dsmwatchd

The **dsmwatchd** command starts the watch daemon.

The watch daemon checks the status of the recall daemon, the monitor daemon, and the scout daemon. If any of these daemons end or become corrupted, **dsmwatchd** automatically recovers the failed daemon.

Valid only on GPFS file systems.

You must have root user authority to use this command.

Restriction: Do not use this command from the console. During installation, it is added to `/etc/inittab`. The **dsmwatchd** command requires a GPFS cluster.

The **dsmwatchd** command writes error messages to the `/dsmerror.log`. If you want **dsmwatchd** to use another `/dsmerror.log` file, you have the following options:

- Use the **errorlogname** option as a command option when you start **dsmwatchd**.
- Use the **errorlogname** option in the `dsm.sys` file to qualify the path and the file name in which to store information about errors that occur during processing. The value of this option overrides the **DSM_LOG** environment variable.
- Set the environment variable, **DSM_LOG**. For example:
`DSM_LOG=/usr/tivoli.tsm/client/hsm/bin/dsmerror.log`
- Restart the **dsmwatchd** daemon. When the **dsmwatchd** daemon starts, it reads the new **errorlogname** option value or the new **DSM_LOG** environment variable setting.

On Red Hat Enterprise Linux Version 6 (RHEL6):

Stop and restart the daemon with the following commands:

```
initctl stop HSM
initctl start HSM
```

On all other systems besides RHEL6:

Stop the **dsmwatchd** daemon by using the command `kill -15 dsmwatchd_pid`. The **dsmwatchd** daemon is restarted by the next **init** process.

You can stop the **dsmwatchd** daemon permanently.

On Red Hat Enterprise Linux Version 6 (RHEL6):

Stop the daemon with the following command: `initctl stop HSM`.

On all other systems besides RHEL6:

Remove the **dsmwatchd** entry from the `/etc/inittab` file.

Note:

1. For AIX GPFS, set the **DSM_LOG** environment variable in the `/etc/environment` file.
2. For Linux GPFS, the `/etc/environment` file does not exist. Use the **errorlogname** option.

Syntax



Parameters

options

-ERRORLOGName=*file_path*

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file `dsm.sys` or `dsm.opt`, or as specified by the **DSM_LOG** environment variable. If any part of the path you specify does not exist, the space management client creates it.

-Help

Displays syntax and options for the command. Do not specify other options when you specify the **help** option.

Chapter 13. HSM GPFS clients Perl scripts

Perl scripts leverage GPFS policy and invoke space management client commands. You can modify the scripts to suit your environment.

The scripts are valid only for HSM General Parallel File System (GPFS) clients.

The scripts assume that you configured GPFS for integration with the space management client.

For information about configuring GPFS integration with the space management client, see Technote 7018848.

dsmMultiServerUpgrade.pl

Use the **dsmMultiServerUpgrade.pl** script to modify space management of a file system so that the file system can be managed by multiple IBM Spectrum Protect servers.

The script couples all files on a file system with the IBM Spectrum Protect server that manages the migration copies and backup versions. The script calls the **dsmreconcile** command to run a special reconciliation that couples the files with the server.

The script is valid only for HSM GPFS clients. The script is located in the `/opt/tivoli/tsm/client/hsm/multiserver/samples/` directory.

Syntax

▶—`dsmMultiServerUpgrade.pl`—`file_system_name`—▶

Parameters

file_system_name

Specify the file system. All files that are migrated, premigrated, archived, or backed up are coupled with the IBM Spectrum Protect server.

Related tasks:

“Enabling a file system to be managed by multiple IBM Spectrum Protect servers” on page 46

dsmNextServername.pl

Use the `dsmNextServername.pl` script to choose the IBM Spectrum Protect server to manage a file. This script is used for automatic migrations.

If a file is selected for migration and is not coupled with any IBM Spectrum Protect server, this script provides a server name. When the file is migrated, it becomes coupled with the IBM Spectrum Protect server.

The script contains settings that are used to choose one IBM Spectrum Protect server from the list of eligible servers. The script is included with the HSM client and by default uses a round-robin method for choosing servers. You can modify

the script to suit your environment. For example, assume that the default server is coupled with many files. When you add a new server, you can exclude the default server until the file management becomes balanced between the new server and the default server. The script is invoked by the GPFS policy engine.

The script is valid only for HSM GPFS clients. The script is located in the `/opt/tivoli/tsm/client/hsm/multiserver/samples/` folder.

Syntax

```
▶▶—dsmNextServername.pl—file_system_name————▶▶
```

Parameters

file_system_name

Specify the file system. Typically, the GPFS policy engine provides the file system name.

dsmreconcileGPFS.pl

Use the **dsmreconcileGPFS.pl** script to reconcile a GPFS file system.

The script uses the GPFS policy engine to determine which files are migrated from the specified file system to IBM Spectrum Protect server storage. The script starts **dsmreconcile** with the **orphancheck** and **fileinfo** options.

The two-way orphan check processing identifies orphans on the file system and on the IBM Spectrum Protect server. The two-way check is done in a single pass and the orphan identification process uses parallel processing.

If a file system is managed by multiple IBM Spectrum Protect servers, the script determines which files are migrated to each IBM Spectrum Protect server.

The script creates a list of files that are coupled with the IBM Spectrum Protect server. The script then starts the **dsmreconcile** command:

```
dsmreconcile file_system_name -servername=server_name -orphancheck -fileinfo  
-filelist=list_file
```

where *listfile* contains a list of all files on the specified file system *file_system_name* that are migrated to the IBM Spectrum Protect server *server_name*.

The script is valid only for HSM GPFS clients. The script is in the `/opt/tivoli/tsm/client/hsm/samples/` directory.

Syntax

```
▶▶—dsmreconcileGPFS—  
└──--tmpdir=temp_dir┐ └──--servername=server_name┐—file_system_name —▶▶
```

Parameters

--tmpdir=temp_dir

Specify a directory for temporary files. By default, the **dsmreconcileGPFS.pl**

script creates temporary files in the /tmp directory. The **dsmreconcileGPFS.pl** script deletes the temporary files when processing is complete. Use two dashes before the **tmpdir** parameter.

-servername=server_name

Specify the IBM Spectrum Protect server. This option is valid only if the file system is managed by multiple IBM Spectrum Protect servers. All files that are migrated to the specified server from the specified file system are reconciled. Obsolete copies on the server are expired and deleted. If the file system is managed by multiple IBM Spectrum Protect servers, you must specify the **servername** option. Use a single dash before the **servername** parameter.

file_system_name

Specify the file system to reconcile. The following tasks are completed for the file system:

- The space management client verifies that valid objects for each local stub file exist in the space management pool of the specified IBM Spectrum Protect server.
- When orphans are identified, their names are recorded in the `.SpaceMan/orphan.stubs` file.

Related reference:

“**dsmreconcile**” on page 161

dsmRemoveServer.pl

Use the **dsmRemoveServer.pl** script to reconcile and recall files that are managed by the specified IBM Spectrum Protect server. The recall and reconciliation processes run for the specified file system.

The `dsmMultiServerRemove.pl` script does the following three tasks:

1. Recalls all migrated files that are coupled with the IBM Spectrum Protect server that was removed to the specified file system.

The space management client determines how much space is required to recall all files to the file system. If there is not enough space, the space management client notifies you. To remove the server, make space available and run the script again.

2. Reconciles the specified file system with the IBM Spectrum Protect server that was removed. The following reconciliation tasks are completed:

- The space management client verifies that valid objects for each local stub file exist in the space management pool of the specified IBM Spectrum Protect server.
- When orphans are identified, their names are recorded in the `.SpaceMan/orphan.stubs` file.
- The space management client expires all files in the space management storage pool on the IBM Spectrum Protect server.

Note: The script does not expire backup copies. If you want to also expire backup copies, you can modify the script by adding the backup-archive client command **dsmc expire**.

The **dsmMultiServerRemove.pl** script calls the `dsmreconcileGPFS.pl` script for the reconciliation.

3. Uncouples all files that are coupled with the IBM Spectrum Protect server that was removed.

More than one recall and reconciliation process can run simultaneously.

The script is valid only for HSM GPFS clients. The script is in the `/opt/tivoli/tsm/client/hsm/multiserver/samples/` directory.

Syntax

```
▶▶—dsmRemoveServer.pl—server=server_name—file_system_name—▶▶
```

Parameters

-server=server_name

Specify the IBM Spectrum Protect server that will no longer manage the file system. Use a single dash before the **server** parameter.

file_system_name

Specify the file system that the IBM Spectrum Protect server will no longer manage.

Related tasks:

“Removing IBM Spectrum Protect servers from a multiple-server environment” on page 48

Related reference:

“dsmreconcileGPFS.pl” on page 170

Chapter 14. Troubleshooting the space management client

Use the following information to resolve some typical problems with the space management client.

Checking the time stamp of a daemon process-ID file

Determine whether a daemon process-ID (PID) file contains a current time stamp. An old time stamp indicates that the daemon stalled.

About this task

HSM daemon processes (**dsmwatchd**, **dsmrecalld**, **dsm scoutd**, and **dsmmonitor**) create a process ID (PID) file for each process. The PID file contains the process ID of the running daemon. The daemon updates the time stamp of the PID file automatically every three seconds. An old time stamp on the PID file for the **dsmwatchd** process or the **dsmrecalld** process indicates a problem with the daemon.

Determine the time stamp of the PID file by completing the following steps:

Procedure

1. List the PID files. Enter the following command:

```
ll /etc/adsm/SpaceMan/*.pid
```

The system displays the process ID files:

```
root@nodeA > ll /etc/adsm/SpaceMan/*.pid
-rw-r--r-- 1 bin bin 6 Nov 09:17 /etc/adsm/SpaceMan/dsmmonitor.pid
-rw-r--r-- 1 bin bin 5 Nov 09:17 /etc/adsm/SpaceMan/dsmrecalld.pid
-rw-r--r-- 1 bin bin 4 Nov 09:17 /etc/adsm/SpaceMan/dsm scoutd.pid
-rw-r--r-- 1 bin bin 3 Nov 09:17 /etc/adsm/SpaceMan/dsmwatchd.pid
```

2. Display the time stamp of the PID file by using the UNIX shell status command and searching for "Change".

Use the command for your operating system:

- On AIX systems, the status command is **istat**.
- On Linux systems, the status command is **stat**.

The following example is for Linux systems:

```
root@nodeA > stat /etc/adsm/SpaceMan/dsmrecalld.pid | grep Change
Change: 2015-11-20 09:26:38.0000000000 +0100
```

What to do next

If the PID file was not updated within the last 3 seconds, start the daemon again.

Related reference:

"dsmmigfs stop, start, and restart" on page 144

Creating a dump file

You can create a dump file without stopping the **dsmwatchd** daemon process.

About this task

You can create a dump file at any time by sending the SIGCONT signal to the **dsmwatchd** daemon.

Restriction: The **dsmwatchd** daemon manages space only on GPFS file systems.

Procedure

1. Query the **dsmwatchd** daemon process ID. Enter the following command:

```
ps -ef | grep dsmwatchd
```

The system returns the process ID of the **dsmwatchd** process.

```
root@nodeA > ps -ef | grep dsmwatchd
root 2068      1 0 Feb08 ?        00:00:44 /opt/tivoli/tsm/client/hsm/bin/dsmwatchd nodetach
root 11766 32142 0 09:14 pts/0    00:00:00 grep --color=auto dsmwatchd
```

2. Send the SIGCONT signal to the **dsmwatchd** process ID by using the **kill** command and the process ID. If the process ID is 2068, enter the following command:

```
kill -SIGCONT 2068
```

Results

The **dsmwatchd** daemon creates two dump files in the `/tmp/hsm/` directory.

1. The `dump.dmapl` file contains current DMAPL session information.
2. The `dump.dsmwatchd` file contains **dsmwatchd** internal counters and **dsmrecalld** daemon information.

IBM Software Support can use the information in the dump files to troubleshoot HSM problems.

HSM entries in the syslog file

The HSM client creates entries in the `syslog` file. The HSM entries in the `syslog` file and the `dsmerror.log` file can identify HSM problems.

The following HSM events are logged to the `syslog` file:

Start and stop of daemon processes

The start time and ID of daemon processes are logged as shown in the following figure.

```
Nov 20 08:48:38 nodeA dsmrecalld: HSM(pid:13756): start master
Nov 20 08:48:38 nodeA dsmrecalld: HSM(pid:13797): start distributor
Nov 20 08:48:38 nodeA dsmrecalld: HSM(pid:13798): start receiver
Nov 20 08:48:38 nodeA dsmrecalld: HSM(pid:24026): start PERMANENT recall worker (ID:3;MIN:3;MAX:5)
Nov 20 08:48:38 nodeA dsmrecalld: HSM(pid:24027): start TEMPORARY recall worker (ID:4;MIN:3;MAX:5)
Nov 20 08:48:38 nodeA dsmrecalld: HSM(pid:24027): stop TEMPORARY recall worker (ID:4)
Nov 20 08:48:38 nodeA dsmrecalld: HSM(pid:20964): stop master
```

Log entries are created for the following daemon types:

Master

The master daemon owns the DMAPI sessions for a space-managed file system. The master daemon responds to data and system events of the DMAPI. Only one master daemon must run on each node where HSM is active. If the master daemon stops, the distributor and receiver daemons also stop.

Distributor

The distributor daemon manages recall operations on the node where the distributor daemon is running. The distributor daemon starts recall worker daemons. Only one distributor daemon must run on each node where HSM is active.

Receiver

The receiver daemon accepts recall requests from other cluster nodes and sends them to the local distributor daemon. Only one receiver daemon must run on each node where HSM is active.

Recall worker

The recall worker daemon completes recall operations.

If there are fewer recall worker daemons than the value of the **MINRECALLDAEMONS** option, the status of the recall worker daemon is permanent (PERMANENT). Otherwise, the status of the recall worker daemon is temporary (TEMPORARY). A temporary daemon is stopped after the file recall operation is finished.

The value of the **MINRECALLDAEMONS** option is indicated by the value of MIN in the log entry. The value of the **MAXRECALLDAEMONS** option is indicated by the value of MAX in the log entry.

Tip: Inspect the recall-worker daemon log entries and determine whether you have to adjust the value of the **MAXRECALLDAEMONS** and **MINRECALLDAEMONS** options to maximize recall processing.

If the syslog file contains many log entries where the recall worker daemon ID value equals the MAX value, increase the value of the **MAXRECALLDAEMONS** option.

If the syslog file indicates that many temporary recall worker daemons are stopped, increase the value of the **MINRECALLDAEMONS** option.

Send and receive signals

All signals that are sent from an HSM process or received by an HSM process are logged as shown in the following figure. Typically, a daemon process stops after the process receives a signal.

```
Nov 20 08:48:09 nodeA dsmwatchd: HSM(pid:7823): signal:15 (Terminated) send to pid:30579
Nov 20 08:48:09 nodeA dsmrecalld: HSM(pid:30579): signal:15 (Terminated) received
```

File system mount events

All mount events that are received by the HSM client are logged as shown in the following figure.

```
Nov 20 08:41:17 nodeA dsmrecalld: HSM(pid:30539): received DM_EVENT_MOUNT for fs:/gpfs2 type:DM_LOCAL_MOUNT
Nov 20 08:41:52 nodeA dsmrecalld: HSM(pid:30539): received DM_EVENT_MOUNT for fs:/gpfs1
```

Creation of dump files

A log entry is created when the **dsmwatchd** daemon creates a dump file.

```
Nov 20 08:41:52 nodeA dsmwatchd: HSM(pid:19418): created dump file: /tmp/hsm/dump.dmapi.2015.12.8.8.43.33
Nov 20 08:41:52 nodeA dsmwatchd: HSM(pid:19418): created dump file: /tmp/hsm/dump.dsmwatchd.2015.12.8.8.43.33
```

System events that stop the recall service

The following conditions can stop the **dsmrecalld** recall service:

- The GPFS file system on a cluster node stops.
- The **dsmrecalld** service does not respond. The PID file time stamp is not updated.
- The number or combination of **dsmrecalld** daemon processes is not correct.

The following log shows system events that stop the recall service.

```
Nov 20 08:40:41 nodeA dsmwatchd: HSM(pid:7823): Stop local recall service. Reason: GPFS down
Nov 20 08:48:32 nodeA dsmwatchd: HSM(pid:7823): Restart local recall service. Reasons: invalid process list
```

Client trace classes for the space management client

The space management client provides individual and aggregate trace classes. You can use aggregate trace classes for a shortcut to enable many related trace classes by specifying the aggregate trace class name.

The trace classes in table Table 22 are used for diagnosing problems with the space management client. To diagnose problems, use the trace class name with the TRACEFLAG options in the dsm.opt file.

Table 22. Trace classes

Trace class	Description	Usage
PID	Displays the process ID on each trace statement. The PID is part of the SERVICE trace aggregate.	Use this trace class to diagnose problems that might involve multiple processes.
SERVICE	Displays general processing information for the client. The SERVICE aggregate includes all trace classes except MEMORY and *DETAIL classes. The SERVICE trace flag can generate a substantial amount of information. Consider using the tracemax option with the SERVICE trace flag.	Use this trace class when the nature of the problem is unknown. If the SERVICE trace flag is used, it is not necessary to specify other trace flags. For HSM daemon tracing, use this flag if the problem can be quickly re-created.
SM	Displays general processing information of the space management client. The SM class is part of the SERVICE or SMALL trace aggregate.	Use this trace class to diagnose problems that are related to HSM (hierarchical storage management).
SMALL	Displays information for HSM.	Use this trace class to enable all trace classes that are related to HSM.

Table 22. Trace classes (continued)

Trace class	Description	Usage
SMXDSM	Displays DMAPI processing information. The SMXDSM class is part of the SERVICE or SMALL trace aggregate.	Use this trace class to diagnose problems that involve DMAPI processing information.
TID	Displays the thread ID on each trace statement. The TID class is part of the SERVICE trace aggregate.	Use this trace class to diagnose problems that might involve multiple threads of a single process.

Related information:

 Troubleshooting: Using trace to resolve problems

Enabling a Space Management client trace

You can trace the space management client by adding trace options to the `dsm.opt` client options file.

Procedure

1. Determine the trace classes that you want to enable for the space management client. See “Client trace classes for the space management client” on page 176 for a list of trace classes.
2. Add trace options to the `dsm.opt` client options file.

- a. Specify the trace classes that you want to enable or disable with the **traceflags** option. You can list one or more trace classes. Separate each trace class with a comma and no space. Use a minus sign (-) in front of a trace class to disable the trace. Put all disabled trace classes at the end of the list. For example, if you want to collect a SERVICE trace without the SESSION or SESSVERB classes, then specify the following option:

```
traceflags service,-session,-sessverb
```

- b. Specify the location of the trace output file by using the **tracefile** option. The **tracefile** option must specify a complete file path. The following example specifies a complete file path:

```
tracefile /home/spike/trace.out
```

- c. Optional: Set a maximum size for the trace output by using the **tracemax** option.

If you specify the **tracemax** option, the trace information wraps and writes over the beginning of the data when the output reaches the specified size. The end of the trace information is indicated by “END OF DATA”. When data wraps, the end of the output contains “Continued at beginning of file”. Consider a maximum trace size if you are trying to capture an event that happens at the end of a long process.

If you specify a **tracemax** value of 1001 or greater and the **tracesegsize** option is not specified, then the trace output is split into several files. Each file is the default segment size, 1000 MB.

Specify the maximum size for the trace output in MB. Valid values are 1 - 4294967295. For example, to set a maximum output size of 4096 MB, use the following option: `tracemax 4096`

- d. Optional: Specify a trace segment size.

Splitting the trace into segments allows easier management of large amounts of trace data. With small files, you do not have to compress or use a separate file splitter utility.

Specify the value in MB. Valid values are 1 - 1000.

Each trace file name contains the value of **tracefile** and appends a number, beginning with 1. For example, if you specify **tracefile** /home/spike/trace.out and **tracesegsize** 200, then the trace is segmented into separate files of no more than 200 MB each. The files are named /home/spike/trace.out.1, /home/spike/trace.out.2, and so on.

3. Run the operation or restart the HSM daemons to enable the daemon tracing.

Results

The space management client creates trace output during the operation.

What to do next

You can also configure and start tracing from a command line. The following example specifies trace options in command format:

```
dsmmigrate -traceflags=service,-session,-sessverb -tracefile=/home/spike/  
trace.out -tracemax=4096 -tracesegsize=200
```

Enabling daemon traces

You can stop and start trace on HSM daemons.

About this task

To trace an HSM daemon, you must add trace options to the **dsm.opt** file and restart the daemon. You must remove the trace options after you restart the daemon, or the space management client traces another daemon when the other daemon restarts. Use different commands and steps to restart different HSM daemons.

Daemon traces can create large trace files in a short time. Enable traces only for selected daemons. Re-create the problem immediately. Disable the daemon trace after the problem occurs.

Important: The HSM daemon holds the file handle for the trace file. Do not delete the trace file when the daemon is running. Deleting a trace file when tracing is enabled can yield unwanted results.

Remember: If the option **hsmdisableautomigdaemons** YES is set in the **dsm.opt** file, the scout (**dsmscoutd**) and monitor (**dsmonitord**) daemons are not running.

Example

Tracing the watch daemon

You can start traces for only the watch daemon (**dsmwatchd**).

Procedure

This procedure relates to a static trace setting.
The setting is applied to all subsequent daemon starts.

1. Stop the HSM daemons (but not the watch daemon) with the following command: `/usr/bin/dsmmigfs stop`.
2. Add trace options to the `dsm.opt` client options file as described in “Enabling a Space Management client trace” on page 177.
3. Stop and restart the watch daemon.
 - On Red Hat Enterprise Linux Version 6 (RHEL6), stop and restart the watch daemon by using the following commands:

```
initctl stop HSM
initctl start HSM
```
 - On all other systems besides RHEL6, stop the watch daemon with the following command:

```
kill -15 dsmwatchd_pid
```

The watch daemon is restarted by the next system **init** process.

4. Remove the trace options from the `dsm.opt` file. If you do not remove the trace options, other daemons are traced when you restart other daemons.
5. Start all HSM daemons other than the watch daemon by using the following command: `/usr/bin/dsmmigfs start` Because the `dsm.opt` file contains no tracing information, the other daemons are not traced.

You can also execute a procedure for a dynamic trace setting.
The setting is applied to the currently running daemon process.

6. Use the command **dsmmigfs control** to enable dynamic tracing. See “**dsmmigfs control**” on page 135 for details.

Results

The space management client captures data from the watch daemon.

What to do next

You can stop tracing the watch daemon in two steps:

1. Verify that the `dsm.opt` client options file does not contain trace options.
2. Stop and restart the watch daemon as described in step 3.

Tracing the recall daemon

You can start traces for only the recall daemon (**dsmrecalld**).

Procedure

This procedure relates to a static trace setting.
The setting is applied to all subsequent daemon starts.

1. In a GPFS environment, issue the following command: `dsmmigfs disableFailover`.
2. Add trace options to the `dsm.opt` client options file as described in “Enabling a Space Management client trace” on page 177.
3. Stop and restart the recall daemon.

- a. Stop the recall daemon by using the following command: **dmkilld**.
 - b. Start the recall daemon by using the following command: **dmrecalld**.
4. Remove the trace options from the `dsm.opt` file. If you do not remove the trace options, other daemons are traced when you restart other daemons.
 5. In a GPFS environment, issue the following command: `dsmmigfs enableFailover`.

You can also execute a procedure for a dynamic trace setting.
The setting is applied to the currently running daemon process.

6. Use the command **dsmmigfs control** to enable dynamic tracing. See “**dsmmigfs control**” on page 135 for details.

Results

The space management client captures data from the recall daemon.

What to do next

You can stop tracing the recall daemon in two steps:

1. Verify that the `dsm.opt` client options file does not contain trace options.
2. Stop and restart the recall daemon as described in step 3 on page 179.

Tracing the scout daemon

You can start traces for only the scout daemon (**dsmscoutd**).

Procedure

1. In a GPFS environment, issue the following command: `dsmmigfs disableFailover`.
2. Add trace options to the `dsm.opt` client options file as described in “Enabling a Space Management client trace” on page 177.
3. Restart the scout daemon by using the following command: **dsmscoutd restart**.
4. Remove the trace options from the `dsm.opt` file. If you do not remove the trace options, other daemons are traced when you restart other daemons.
5. In a GPFS environment, issue the following command: `dsmmigfs enableFailover`.

Results

The space management client captures data from the scout daemon.

What to do next

You can stop tracing the scout daemon in two steps:

1. Verify that the `dsm.opt` client options file does not contain trace options.
2. Restart the scout daemon by using the **dsmscoutd restart** command.

Tracing the monitor daemon

You can start traces for only the monitor daemon (**dsmonitord**).

Procedure

1. In a GPFS environment, issue the following command: `dsmmigfs disableFailover`.
2. Add trace options to the `dsm.opt` client options file as described in “Enabling a Space Management client trace” on page 177.
3. Stop and restart the monitor daemon.
 - a. Stop the monitor daemon by using the following command: `kill -15 dsmonitord_pid`.
 - b. Start the monitor daemon by using the following command: **dsmonitord**.
4. Remove the trace options from the `dsm.opt` file. If you do not remove the trace options, other daemons are traced when you restart other daemons.
5. In a GPFS environment, issue the following command: `dsmmigfs enableFailover`.

Results

The space management client captures data from the monitor daemon.

What to do next

You can stop tracing the monitor daemon in two steps:

1. Verify that the `dsm.opt` client options file does not contain trace options.
2. Stop and restart the monitor daemon as described in 3.

Log files to aid troubleshooting

When troubleshooting the space management client, you can configure the logs that record HSM activity and error messages.

Related reference:

“Logs for HSM activity and error messages” on page 32

Related information:

 Backup-archive client options reference

Commands for displaying HSM information

space management client commands are provided to display HSM statistics and configuration information.

The following commands yield HSM information for troubleshooting. The commands are organized by the type of information that they yield.

Cluster-specific information

Enter these commands on one node in the cluster.

- To display the configuration and state information for the space management file system, enter the following command:
`/usr/bin/ls -aLR /file_system_name/.SpaceMan`
- To display the space management file system and cluster configuration, enter the following command:

```
/usr/bin/dsmmigfs query -detail
```

- To display the space management failover configuration, enter the following command:

```
/usr/bin/dsmmigfs query -failover
```

- To display the configured IBM Spectrum Protect options, enter the following command:

```
/usr/bin/dsmmigquery -options
```

- To display details on the assigned IBM Spectrum Protect server management class configuration, enter the following command:

```
/usr/bin/dsmmigquery -mgmt -detail
```

Node-specific information

Enter these commands on all nodes in the cluster.

- To display the content of the node-specific space management configuration, enter the following command:

```
/usr/bin/ls -aLR /etc/adsm/SpaceMan
```

- To display the client option settings, enter the following command:

On AIX systems:

```
/usr/bin/cat /usr/tivoli/tsm/client/ba/bin64/dsm.opt
```

On Linux systems:

```
/usr/bin/cat /opt/tivoli/tsm/client/ba/bin/dsm.opt
```

- To display the administrator option settings, enter the following command:

On AIX systems:

```
/usr/bin/cat /usr/tivoli/tsm/client/ba/bin64/dsm.sys
```

On Linux systems:

```
/usr/bin/cat /opt/tivoli/tsm/client/ba/bin/dsm.sys
```

- To display the installed space management client executable programs, enter the following command:

On AIX systems:

```
/usr/bin/ls -aLR /usr/tivoli/tsm/client/hsm/bin
```

On Linux systems:

```
/usr/bin/ls -aLR /opt/tivoli/tsm/client/hsm/bin
```

Process-specific information

Enter these commands to display process information.

- To display all process IDs, enter the following command:

```
/usr/bin/ps -ef |grep dsm
```

- To display all process stack information, enter the following command:

On AIX systems:

```
/usr/bin/procstack process_ID
```

On Linux systems, use one of the following commands:

```
/usr/bin/pstack process_ID
```

or

```
/usr/bin/gstack process_ID
```

Commands for displaying GPFS information

General Parallel File System (GPFS) commands are provided to display General Parallel File System (GPFS) information.

The following commands yield GPFS information for troubleshooting. The commands are organized by the type of information that they yield.

Cluster-specific information

Enter these commands on one node in the cluster.

- To display the configuration information for a GPFS cluster, enter the following command:

```
/usr/lpp/mmfs/bin/mmlscluster
```

- To display the available file space on a GPFS file system, enter the following command:

```
/usr/lpp/mmfs/bin/mmdf device
```

- To display all the GPFS policy information for a specified file system, enter the following command:

```
/usr/lpp/mmfs/bin/mmlspolicy device -L
```

- To display the state of the GPFS daemon on all nodes in the cluster, enter the following command:

```
usr/lpp/mmfs/bin/mmgetstate -a
```

- To display the cluster configuration data for a GPFS cluster, enter the following command:

```
/usr/lpp/mmfs/bin/mmlsconfig
```

- To create a memory dump of the GPFS DMAPI information, enter the following command:

```
/usr/lpp/mmfs/bin/mmfsadm dump dmapi
```

Node-specific information

To display the GPFS node configuration, enter the following command on all nodes in the cluster:

```
/usr/bin/cat /var/mmfs/gen/mmsdrfs
```

For information about GPFS commands and GPFS requirements for the IBM Spectrum Protect for Space Management client, go to the General Parallel File Systems product information and see `mmbackup` command: IBM Spectrum Protect requirements.

Commands for displaying operating system information

Operating system commands are provided to display operating system information.

To display the version and release level of the operating system, enter the following command:

On AIX systems:

```
oslevel -r
```

On Linux systems:

```
uname -a
```

Other tools for collecting information

A guide to collecting information is available from IBMSoftware Support.

For more information about gathering data to troubleshoot the space management client, see Technote 1268553.

The document contains a Perl script that is named `collect_v8.pl`, which enables HSM information to be collected automatically.

HSM common problems and solutions

space management client common problems are listed. Typical solutions are suggested.

The following table lists common problems and typical solutions.

Table 23. Common HSM problems and resolutions

Problem	Problem source	Solution
No HSM daemons are running.	The configuration in the <code>dsm.opt</code> file or the <code>dsm.sys</code> file is invalid. The error prevents all HSM daemons from starting.	Run any HSM command. The command output describes the failure. Correct the configuration in the <code>dsm.opt</code> file or the <code>dsm.sys</code> file.
The watch daemon (dsmwatchd) is the only active daemon.	Any of the following conditions can cause this problem: <ul style="list-style-type: none"> HSM was stopped on the specified node. Failover is disabled on the specified node. The DMAPi service is not running. 	Try the following solutions: <ul style="list-style-type: none"> Start the HSM daemons by issuing the HSM command: dsmmigfs start. The daemons might take up to 30 seconds to start running. Enable the failover on the node by issuing the HSM command: dsmmigfs enablefailover Ensure that GPFS is in the active state on all nodes in the cluster. To verify this state, issue the GPFS command: mmgetstate -a
The mount of DMAPi-enabled file systems fails.	The recall daemon does not run.	Ensure that the recall daemon runs by issuing the command: dsmrecalld . The mount of a DMAPi-enabled file system requires at least one recall daemon in the cluster to be running.
The mount of DMAPi-enabled file systems hangs	There are two possible causes: <ol style="list-style-type: none"> On one node in the GPFS cluster, there is an orphaned DMAPi session from a recall daemon that failed. The GPFS configuration parameter, enableLowSpaceEvents, is set to yes. To see the current value of this parameter, issue the command: <pre>mmfsconfig grep enableLowSpaceEvents</pre> 	If there is an orphaned DMAPi session, restart the recall daemon: <ol style="list-style-type: none"> Stop the recall daemon on all the nodes in the cluster. Issue the command: dmkilld. Start the recall daemon by issuing the command: dsmrecalld. The orphaned DMAPi session is cleaned up while the recall daemon starts. <p>If enableLowSpaceEvents=yes, change the value and restart the GPFS daemon on all nodes:</p> <ol style="list-style-type: none"> <code>/usr/lpp/mmfs/bin/mmchconfig enableLowSpaceEvents=no</code> <code>/usr/lpp/mmfs/bin/mmshutdown -a</code> <code>/usr/lpp/mmfs/bin/mmstartup -a</code>
Several Space Management commands end without processing.	The space management client cannot access the node configuration in the <code>/etc/adsm/SpaceMan</code> directory. Typically, this condition is caused by an unmounted <code>/etc</code> file system.	Mount the <code>/etc</code> file system.

Table 23. Common HSM problems and resolutions (continued)

Problem	Problem source	Solution
<p>A file migration operation fails and displays the messages "ANS1228E Sending of object .. failed." and "ANS9256E File .. is currently opened by another process." or a file recall operation hangs and is unable to provide user feedback.</p>	<p>A previous file migration operation or file recall operation of the affected file terminated prematurely. This termination was caused by a failure or a GPFS shutdown on the node that processed the file migration operation or file recall operation. Later on this node either:</p> <ol style="list-style-type: none"> 1. The recall daemons were restarted BEFORE the affected file system was remounted, or 2. The recall daemons were not restarted at all, or 3. The recall daemons were restarted on this node even though this node was not the owner of the affected file system. 	<p>Restart the recall daemons on the affected node by issuing the command: dsmmigfs restart. If it is not clear which node caused this problem, perform the following procedure:</p> <ol style="list-style-type: none"> 1. Recursively list the content of the <code>.SpaceMan/logdir/</code> subdirectory of the affected file system, for example, <code>ls -lR /<affected_file_system>/.SpaceMan/logdir/</code>. 2. Look for entries in all <code>translog</code> subdirectories that contain one or more trailing digits in their name, for example, <code>translog12/</code>. <ul style="list-style-type: none"> • For problems with a file migration operation, look for entries with the suffix <code>.mig</code>, for example, <code>099B3477562F877D0000000000007D5B000000000000001772E70200000000.mig</code> • For problems with a file recall operation, look for entries with suffix <code>.rec</code>, for example, <code>099B3477562F877D0000000000007D5B000000000000001772E70200000000.rec</code> 3. Take note of the trailing digits in the name of the matching <code>translog</code> subdirectory or subdirectories, for example, "12" for the subdirectory <code>translog12/</code>. 4. Issue the command: mmfscluster on one of the cluster nodes. In the command's output, look for a number in the "Node" column that matches the trailing digits. For example, look for "12" in this column. In the same row, look for the node name that corresponds with the matching number, for example, "12" corresponds with node name "number_cruncher". 5. Ensure that the matching node is the owner of the affected file system. Check for the owner by issuing the command: dsmmigfs q -d. If required, transfer the file system ownership to the node identified in step 4 by issuing the command: dsmmigfs takeover /<affected_file_system> on this node. 6. Restart the recall daemons on the matching nodes, for example, on the node "number_cruncher". <p>Important: When you restart the recall daemons, make sure not to disrupt unaffected recall processes.</p>

Appendix A. Control files in the .SpaceMan directory

When you add space management to your file systems, the space management client creates control files in a hidden directory named `.SpaceMan`.

Table 24 provides a brief description of the control files and directories in the `.SpaceMan` directory in each space-managed file system. These files are also in the `/etc/adsm/SpaceMan` directory that is created when you install the space management client.

The HSM client automatically excludes these files from space management. No action on your part is required to ensure that the files remain on the local file systems. The space management client requires these objects for processing. Do not delete or modify the `.SpaceMan` directory, its contents, or any of the file ownerships and permissions.

Table 24. Control files stored in the .SpaceMan directory

Files	Description
<code>ActiveRecallTab</code>	This file contains the active recall table.
<code>config/dmiFSGlobalState</code>	This file contains global and file system information about the Data Management Application Programming Interface (DMAPI) interface.
<code>config/dmiFSGlobalState.pid</code>	This file is a lock file for the <code>dmiFSGlobalState</code> file.
<code>dmiFSState</code>	This file stores information about the file system.
<code>dsmmigfstab</code>	If this file exists, it is from an older version of the HSM client. It is needed for transition to current version.
<code>hsmfsconfig.pid</code>	This file is a lock file for the local <code>hsmfsconfig.xml</code> file.
<code>hsmfsconfig.xml</code>	This file contains the space management settings for the file system.
<code>logdir/</code>	This directory records information during file migration or recall processing. The information is used to complete any interrupted transactions, such as during a system failure.
<code>metadata/</code>	This directory contains the complete file index (CFI) that was created by the scout daemon. The CFI contains file system information.
<code>orphan.stubs</code>	This file records the orphan files that were identified during reconciliation.
<code>reserved/</code>	This directory contains reserved files for demand migration.
<code>multiserver/BasicRuleSet</code>	This file is a rule set for a file system that is managed by multiple IBM Spectrum Protect servers.
<code>multiserver/serverlist</code>	This file contains the list of IBM Spectrum Protect servers that manage the file system.
<code>ruleset</code>	This file contains the automatically generated rule set for GPFS.
<code>SDR/</code>	If this directory exists, it is from an older version of the space management client.

Table 24. Control files stored in the .SpaceMan directory (continued)

Files	Description
status	This file records space management statistics for the file system.
tapeOptimizedRecall/	This directory contains the recall list files that optimize tape access.

For HSM storage pool support, all files located in the .SpaceMan directory are placed in the same storage pool. EXCLUDE rules prevent other migration rules from moving those files to a different pool.

To ensure that other migration rules do not move files into the .SpaceMan directory, you can specify the following rules:

```
RULE 'TSM_EXCL_DOTSPACEMAN' EXCLUDE WHERE PATH_NAME LIKE '%/.SpaceMan/%'
```

If files were moved to an unwanted location, the following rule moves all files into one storage pool:

```
RULE 'TSM_MIGR_DOTSPACEMAN' MIGRATE TO POOL 'target_pool'
WHERE PATH_NAME LIKE '%/.SpaceMan/%'
```

where *target_pool* might be "system" or the custom default pool if the system pool is dedicated only to metadata.

Related concepts:

Chapter 8, "File system reconciliation," on page 79

"Managing a file system with multiple IBM Spectrum Protect servers" on page 45

"List files for optimized tape recalls" on page 76

Space requirement for HSM control files

The space management client creates control files which require free space on the system. Without sufficient free space, HSM cannot run properly.

Most files in the .SpaceMan directory occupy approximately 1 GB of space. In addition to this basic requirement, there are some files that can occupy significantly more space.

.SpaceMan/metadata

If the scout daemon is running, the daemon creates the complete file index (CFI). The scout daemon runs when you set **hsmdisableautomigdaemons=no** (the default value). The CFI requires approximately 1 KB per file system block or file. For example, if the file system size is 1 TB, and the block size is 1 MB, there are 1,000,000 blocks on the file system. Multiply 1,000,000 blocks by 1 KB per block to yield a CFI size of 1,000,000,000 (1 GB).

You can set the size of the CFI by specifying a maximum number of files to be space managed. Use the **maxfiles** option with the **dsmmigfs update** command. For example, if the file system block size is 1 KB and you specify **maxfiles=1000**, the CFI size is 1000 × 1 KB = 1 MB. In this example the CFI is only 1 MB, regardless of the total size of the file system. The value of the **maxfiles** option must be greater than or equal to the number of files in the file system.

.SpaceMan/reserved

When you add space management to a file system, the space management client creates reserved files. The reserved files require 1 GB of space for each file system that is space managed.

The total space required can be estimated as 1 MB for base information plus 1 GB for reserved files for each space-managed file system. If **hsmdisableautomigdaemons=no** (the default value), you must also add space for the CFI.

Related reference:

“dsmmigfs add and update” on page 128

Appendix B. Accessibility features for the IBM Spectrum Protect product family

Accessibility features assist users who have a disability, such as restricted mobility or limited vision, to use information technology content successfully.

Overview

The IBM Spectrum Protect family of products includes the following major accessibility features:

- Keyboard-only operation
- Operations that use a screen reader

The IBM Spectrum Protect family of products uses the latest W3C Standard, WAI-ARIA 1.0 (www.w3.org/TR/wai-aria/), to ensure compliance with US Section 508 (www.access-board.gov/guidelines-and-standards/communications-and-it/about-the-section-508-standards/section-508-standards) and Web Content Accessibility Guidelines (WCAG) 2.0 (www.w3.org/TR/WCAG20/). To take advantage of accessibility features, use the latest release of your screen reader and the latest web browser that is supported by the product.

The product documentation in IBM Knowledge Center is enabled for accessibility. The accessibility features of IBM Knowledge Center are described in the Accessibility section of the IBM Knowledge Center help (www.ibm.com/support/knowledgecenter/about/releasenotes.html?view=kc#accessibility).

Keyboard navigation

This product uses standard navigation keys.

Interface information

User interfaces do not have content that flashes 2 - 55 times per second.

Web user interfaces rely on cascading style sheets to render content properly and to provide a usable experience. The application provides an equivalent way for low-vision users to use system display settings, including high-contrast mode. You can control font size by using the device or web browser settings.

Web user interfaces include WAI-ARIA navigational landmarks that you can use to quickly navigate to functional areas in the application.

Vendor software

The IBM Spectrum Protect product family includes certain vendor software that is not covered under the IBM license agreement. IBM makes no representation about the accessibility features of these products. Contact the vendor for accessibility information about its products.

Related accessibility information

In addition to standard IBM help desk and support websites, IBM has a TTY telephone service for use by deaf or hard of hearing customers to access sales and support services:

TTY service
800-IBM-3383 (800-426-3383)
(within North America)

For more information about the commitment that IBM has to accessibility, see IBM Accessibility (www.ibm.com/able).

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. 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 IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

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

*IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan*

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those

websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

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 Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US*

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

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows: © (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. _enter the year or years_.

Trademarks

IBM, the IBM logo, and ibm.com[®] are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Adobe is a registered trademark of Adobe Systems Incorporated in the United States, and/or other countries.

Linear Tape-Open, LTO, and Ultrium are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

Intel and Itanium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java[™] and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

SoftLayer[®] is a registered trademark of SoftLayer, Inc., an IBM Company.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Privacy policy considerations

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at <http://www.ibm.com/privacy> and IBM's Online Privacy Statement at <http://www.ibm.com/privacy/details> in the section entitled "Cookies, Web Beacons and Other Technologies," and the "IBM Software Products and Software-as-a-Service Privacy Statement" at <http://www.ibm.com/software/info/product-privacy>.

Glossary

A glossary is available with terms and definitions for the IBM Spectrum Protect family of products.

See the IBM Spectrum Protect glossary.

To view glossaries for other IBM products, see IBM Terminology.

Index

Special characters

.SpaceMan directory 40, 187

.SpaceMan/orphan.stubs 91

A

accessibility features 191

activating space management 137

add 128

adding HSM

 dsmmigfs command 128

 exported file systems 37

 multiple servers 47, 134

 nested file systems 36

 overview 35

 to file system 36

 WPARs 39

afmskipuncachedfiles 89

AIX

 installing HSM 9

AIX GPFS

 installing HSM 7, 9

 upgrading 10

 uninstalling HSM 11

archive command 63

archiving

 file systems with HSM 61, 62

archiving migrated files 63

automatic migration 53

 by GPFS policy 2, 53, 55

 dsmNextServername.pl 169

 file prioritization 55

 options 56

 threshold migration 121

 start manually 57

automignonuse 25

B

backing up

 before migrating 63

 file systems with HSM 61

backing up migrated files 63

backup image command 63

backup-archive client 1

Bourne and Korn shell 29

Bourne shell variables 29

C

C shell 29

C shell variables 29

candidate selection

 migration 56

candidates file 187

candidatesinterval 90

checkfororphans 90

checking daemons 173

checkthresholds 91

clusters on GPFS

 adding space management 31

command

 dsmmighelp 114

command line

 return codes for operations 115

command line help

 displaying 114

commands

 archive 63

 backup image 63

 descriptions 116

 displaying GPFS information 183

 displaying HSM information 181

 displaying operating system information 183

 dmkilld 118

 dsmattr 118

 dsmautomig 121

 start threshold migration 57

 dsmdf 114, 123

 young files are omitted 95

 dsmdu 114, 123, 124

 young files are omitted 95

 dsmls 114, 125, 126, 157

 dsmmigfs 114, 139, 144

 adding HSM 36

 adding or updating 128

 deactivating space management 51, 137

 globally deactivating space management 138

 help 139

 reactivating space management 51, 137

 removing space management 137

 takeover 146

 dsmmigfs ADDMultiserver 134

 dsmmigfs command

 globally deactivating space management 138

 globally reactivating space management 138

 dsmmigfs control 135

 dsmmigfs QUERYMultiserver 134

 dsmmigfs REMOVEMultiserver 134

 dsmmigfs rollback 142

 dsmmighelp 147

 dsmmigquery 26, 147

 dsmmigrate 149

 selective migration 58

 dsmmigundelete 114, 153

 restoring stub files 65

 dsmonitord 155

 space monitor daemon 83

 dsmq 156

 dsmrecall 157

 dsmrecalld 160, 166

 recall daemon 84

 dsmreconcile 161

 dsmrm 163

 dsmscoutd 164

 dsmscoutd command 166

 dsmsetpw 165

 dsmwatchd 166

 watch daemon 85

 incremental 63

- commands (*continued*)
 - selective 63
 - using 113
 - using in executables 115
 - using in shell scripts 115
- compressalways 92
- compression 92
- configuring HSM
 - overview 35
- configuring HSM connection
 - to secondary server 23
- configuring space management client 19
- control files
 - HSM 187
- creating stub files for premigrated files 153

D

- daemon
 - dsmmonitor 122
 - dsm scout 43
- daemons
 - dsmmonitor 155, 161
 - dsmrecall 160
 - dsm scout 164
 - dsmwatchd 166
 - dsmwatchd daemon 35
 - overview 83
 - recall daemon 84
 - scout daemon 84
 - space monitor daemon 83
 - stopping 85
 - watch daemon 85
- deactivating HSM 51
- deactivating space management 137
 - globally 138
- defaultserver 22, 92
- defined 3
- deleted stub files, restoring 153
- demand migration 2, 53, 55
- different backup and migration servers 68
- disability 191
- disable failover 143
- displaying
 - space management information 124
 - space management settings 139
 - space usage information 124
- displaying GPFS information 183
- displaying HSM information 181
- displaying online command help 114
- displaying operating system information 183
- displaying option information 20
- DMAPI 35
- dmkilld command 118
- DSM_CONFIG 28
- DSM_CONFIG environment variable 28
- DSM_DIR environment variable 28
- DSM_LOG environment variable 28
- dsm.opt
 - editing 20, 22
- dsm.sys
 - editing 20
 - options 20
- dsmattr command 118
- dsmautomig command 121
 - start threshold migration 57
- dsmddf command 114, 123

- dsmdu command 114, 123, 124
- dsmerror.log 29
- dsmls command 114, 125, 126, 157
- dsmmigfs
 - command 146
 - deactivating space management 51
 - reactivating space management 51
 - control 135
 - deactivating space management 137
 - GPFS 134
 - help 139
 - query 139
 - reactivating space management 137
- dsmmigfs ADDMultiserver 134
- dsmmigfs command 114, 142, 144
 - adding HSM 36
 - adding space management 128
 - updating space management settings 128
- dsmmigfs control 135
- dsmmigfs DISABLEFailover 143
- dsmmigfs ENABLEFailover 143
- dsmmigfs QUERYMultiserver 134
- dsmmigfs REMOVEMultiserver 134
- dsmmighelp command 114, 147
- dsmmigquery
 - shared options 149
- dsmmigquery command 26, 147
- dsmmigrate command 149
 - selective migration 58
 - start selective migration 58
- dsmmigundelete command 114, 153
 - restoring stub files 66
- dsmmonitor command 155
 - space monitor daemon 83
- dsmmonitor daemon 122, 161
- dsmMultiServerUpgrade.pl 46, 169
- dsmNextServername.pl 169
- dsmq command 156
- dsmrecall command 157
- dsmrecalld command 160, 166
 - recall daemon 84
- dsmreconcile
 - reconcileinterval 4
- dsmreconcile command 161
- dsmreconcileGPFS.pl 170
- dsmRemoveServer.pl 171
- dsmrm command 163
- dsm scoutd command 164, 166
- dsmsetpw command 165
- dsmwatchd command 166
 - watch daemon 85
- dump files 174

E

- editing
 - dsm.opt 20
 - dsm.sys 20
- enable failover 143
- enablelanfree 30
- encrypted file system 128
- enhancements for V8.1.2 xiii
- environment variables
 - DSM_CONFIG 28
 - DSM_DIR 28
 - DSM_LOG 28
- error log 32

- error messages 32
- errorlogname 93
- errorprog 94
- exclude
 - files (LAN-free) 30
- exclude list
 - edit 28
- exclude options 26
- exclude.compression 92
- excluding files
 - from backup 26
 - from migration 26
 - from space management 26
- executable file
 - return codes from 115
- expiration 4
- exported file systems
 - adding HSM 37

F

- failover configuration 23, 50
- file list 126
- file migration 53
- file systems
 - backing up 61
 - reconcile 4
 - reconciling
 - automatic reconciliation 79
 - GPFS policy 82
 - immediate reconciliation 81
 - manual reconciliation 80
 - orphan stub files 81
 - overview 79
 - restoring 61
 - different backup and migration servers 68
 - from multiple servers 69
 - same backup and migration server 68
- filelist parameter 157
- files
 - archiving 61
 - backing up 61
 - migrating 61
 - restoring 61
 - retrieving 61
- free space requirement
 - HSM control files 188
- from multiple servers 69

G

- globally deactivating space management 138
- globally reactivating space management 138
- GPFS
 - AFM support 32
 - failover 143
 - GPFS policy reconciliation 82
 - limitations 32
 - policy-driven automatic migration 53, 55
 - re-configuring the hostname or GPFS node number 37

H

- help 114
- high threshold
 - setting 128

- high threshold for migration 42
- HSM 187
- HSM activity log 32
- HSM control files 187
 - space requirement 188
- HSM error messages log 32
- HSM GPFS clients
 - scripts 169
- HSM log 32
- HSM settings
 - high and low thresholds 42
 - maximum number of files 45
 - migration quotas 43
 - minimum file size for migration 42
 - minimum partial file size 45
 - minimum streaming file size 44
 - overview 40
 - premigration 43
 - stub file size 44
 - updating
 - command line 40
- hsmdisableautomigdaemons 46, 94
- hsmdistributedrecall 95
- hsmenableimmediatemigrate 95
- hsmenableimmediatemigrate option 54
- hsmeventdestroy 96
- hsmextobjidattr 97
- hsmfsconfig.xml 35
- hsmfsconfig.xml file 40
- hsmgroupedmigrate 22, 97
- hsmlogeventflags 98
- hsmlogmax 99
- hsmlogname 99
- hsmlogretention 100
- hsmlogsampleinterval 101
- hsmmaxrecalltapedrives 101
- hsmmigzeroblockfiles 102
- hsmmultiserver 46, 103
- hsthreshold 128

I

- IBM Knowledge Center xi
- IBM Spectrum Protect password 165
- immediate reconciliation 81
- inclexcl 103
- inclexcl option
 - dsm.sys 26
- include
 - files (LAN-free) 30
- include list
 - edit 28
- include options 26
- include-exclude
 - file 19
- include.compression 92
- including files
 - for backup 26
 - for migration 26
- incremental command 63
- inodes 123
- installing HSM 5
 - AIX 9
 - AIX GPFS 7, 9
 - installation directory 5
 - Linux GPFS 12, 14
 - prerequisites 5

- installing HSM (*continued*)
 - requirements 5
 - Linux GPFS 14
 - uninstalling
 - AIX GPFS 11
 - Linux GPFS 16
 - upgrade 5
 - upgrading
 - AIX GPFS 10
 - Linux GPFS 15

K

- keyboard 191
- Knowledge Center xi
- Korn shell variables 29

L

- LAN-free 27, 30
 - prerequisites 30
- LAN-free configuration
 - limitations 49
- LAN-free data transfer 29
- lanfreecommmethod 30
- lanfreetcpport 30
- limitations
 - GPFS 32
 - HSM 5
 - LAN-free configuration 49
 - multiple servers 49
 - server name encryption 49
 - server node replication 49
- Linux GPFS
 - installing HSM 12, 14
 - upgrading 15
 - requirements
 - installing HSM 14
 - uninstalling HSM 16
- list file
 - format 157
- list of files
 - displaying 147
- log
 - HSM activity 32
 - HSM error messages 32
- log files
 - dsmerror.log 29
- logdir directory 187
- low threshold
 - setting 128
- low threshold for migration 42
- lthreshold 128

M

- manage recovery 143
- management class
 - assign to files 24
 - default 24
 - displaying 26
- manual reconciliation 80
- maxcandidates 128
- maxcandprocs 104
- maxfiles 128
- maximum number of files 45

- maxmigrators 104
- maxrecalldaemons 105
- maxthresholdproc 105
- messages 32
- migdestination 25
- migfileexpiration 106
- migrated file
 - setting a recall mode 118
- migrateserver 22, 106
- migrating
 - demand 2
 - overview 2
 - recalling files 3
 - selective 2
 - threshold 2
- migrating files
 - automatically 121
 - selectively 149
- migration 42
 - automatic 53, 55
 - dsmNextServername.pl 169
 - options 56
 - by GPFS policy 53
 - candidate selection 56
 - demand 53, 55
 - eligibility criteria 54
 - high and low thresholds 42
 - multiple servers
 - overview 59
 - overview 53
 - premigration 55
 - selective 53, 58
 - selective migration 149
 - start manually 58
 - threshold 53, 55
 - threshold migration 121
 - start manually 57
- migration candidates list 161
 - displaying 147
- migration quotas 43
- migration types 53
- migrequiresbkup 25
- minimum file size 42
- minimum file size for migration 42
- minmigfilesize 107
- minpartialrecallsz
 - setting 128
- minrecalldaemons 108
- minstreamfilesize 128
- mount point
 - changing the mount point of a space-managed file system
 - renaming the file space after mounting at the new mount point 38
 - renaming the file space before mounting at the new mount point 38
- multiple server
 - dsmreconcileGPFS.pl 170
- multiple servers
 - adding HSM 47
 - dsmMultiServerUpgrade.pl 169
 - dsmNextServername.pl 169
 - dsmRemoveServer.pl 171
 - enabling HSM 46
 - limitations 49
 - node replication 50
 - overview 45
 - reconciliation 82

- multiple servers (*continued*)
 - removing a server 48
 - restoring file systems 69
 - scripts 169

N

- nested file systems
 - adding HSM 36
- NFS servers 37
- no-query restore 65
- node replication
 - configuration 23
 - multiple servers 50
- normal recall 72
 - setting for a migrated file 118

O

- obsolete files
 - reconciliation 4
- online help
 - displaying for commands 114
- optimized tape recall 3
 - list files 76
 - overview 75
 - procedure 78
- optional setup 24
- options 30
 - afmskipuncachedfiles 89
 - candidatesinterval 90
 - checkfororphans 90
 - checkthresholds 91
 - compressalways 92
 - compression 92
 - defaultserver 22, 92
 - displaying 20
 - dsm.opt 89
 - dsm.sys 89
 - enablelanfree 30
 - errorlogname 93
 - errorprog 94
 - exclude 26
 - exclude.compression 92
 - hsmdisableautomigdaemons 94
 - hsmdistributedrecall 95
 - hsmenableimmediatemigrate 54, 95
 - hsmeventdestroy 96
 - hsmextobjidattr 97
 - hsmgroupedmigrate 22, 97
 - hsmlogeventflags 98
 - hsmlogmax 99
 - hsmlogname 99
 - hsmlogretention 100
 - hsmlogsampleinterval 101
 - hsmmaxrecalltapedrives 101
 - hsmmigrzeroblockfiles 102
 - hsmmultiserver 103
 - includ 103
 - include 26
 - include.compression 92
 - lanfreecommmethod 30
 - lanfreetcpport 30
 - maxcandprocs 104
 - maxmigrators 104
 - maxrecalldaemons 105

- options (*continued*)
 - maxthresholdproc 105
 - migfileexpiration 106
 - migrateserver 22, 106
 - minmigfilesize 107
 - minrecalldaemons 108
 - passwordaccess 18
 - preservelastaccessdate 54
 - reconcileinterval 108
 - restoremigstate 22, 109
 - rollback 142
 - setting space management 19
 - skipmigrated 63, 110
 - standard format 113
- ordered recall list
 - displaying 147
- orphan stubs 80
- orphan.stubs file 187
- overview 5, 40

P

- partial file recall 3, 72
 - minimum partial file size 45
- partial recall mode
 - setting for a migrated file 118
- password 165
 - encryption 18
 - generate 18
 - managing 18
 - passwordaccess 18
- passwordaccess option 18
- perl scripts 169
 - dsmMultiServerUpgrade.pl 169
 - dsmNextServername.pl 169
 - dsmreconcileGPFS.pl 170
 - dsmRemoveServer.pl 171
- pmpercentage 128
- policies
 - displaying 24, 26
- premigration 55
- premigration percentage 43, 55
- premigration setting 128
- premgrdb.dir file 187
- premgrdb.pag file 187
- prerequisites 29
 - installing HSM 5
 - LAN-free 30
- preservelastaccessdate option 54
- process ID file 173
- publications xi

Q

- quota 128
 - setting 128
- quotas 43

R

- reactivating HSM 51
- reactivating space management 137
 - globally 138
- recall daemon 84
 - dmkilld command 118
 - starting 160

- recall ID 156
- recall mode
 - normal 118
 - setting for a migrated file 118
- recall process
 - displaying status of 156
- recalling
 - mode selection 74
 - normal recall 72
 - overview 71
 - partial file recall 72
 - selective recall 72
 - setting the recall mode 74
 - streaming recall 73
 - tape-optimized recall
 - list files 76
 - overview 75
 - procedure 78
 - transparent recall 71
- recalling files
 - canceling 163
 - migrated files 3
 - removing from queue 163
 - selectively 157
- reconcile file systems 4
- reconcileinterval 108
- reconciliation
 - dsmreconcileGPFS.pl 170
 - policy-based 161
- reconciling
 - automatic reconciliation 79
 - GPFS policy 82
 - immediate reconciliation 81
 - manual reconciliation 80
 - manually 161
 - orphan stub files 81
 - overview 79
- reconciling file systems 161
- reconfiguring HSM
 - GPFS clusters 37
- recreating deleted stub files 153
- registration
 - HSM client 17
 - closed registration 18
 - open registration 17
- remounting a space-managed file system 38
- removing
 - space management 137
- removing a server 48
- removing HSM 52
 - GPFS clusters 37
- requirements
 - installing HSM 5
 - AIX GPFS 9
 - Linux GPFS 14
- resolving 80
- restarting space management daemons 144
- restoremigstate 22, 109
- restoring 68, 69
 - deleted stub files 153
 - dsmmigundelete command 66
 - file systems with HSM 61
 - no-query restore 65
 - standard-query restore 65
 - stub files 66
- restoring file systems
 - overview 67

- restrictions
 - HSM 5
- retrieving
 - file systems with HSM 63
- retrieving files 61
- return codes for operations 115
- rollback option 142
- root user tasks 155
 - creating stub files for premigrated files 153
 - recreating deleted stub files 153
 - starting a scout daemon 164

S

- same backup and migration server 68
- SAN
 - LAN-free 29
- scheduling
 - options 87
 - overview 87
 - starting 87
- scout daemon 84
 - starting 164
- scripts
 - for HSM GPFS clients 169
 - perl 169
- selected files
 - migrating 149
 - recalling 157
- selecting HSM options 19
- selective command 63
- selective migration 2, 53, 58
 - performing 149
 - start migration 58
- selective recall 3, 72, 157
- server name encryption
 - limitations 49
- server node replication
 - limitations 49
- services
 - scheduling
 - options 87
 - overview 87
 - starting 87
- setting environment variables 28
- setting space management options 19
- setting the recall mode 74
- setting up 19
- setup tasks 24
- shell scripts 1
 - return codes from 115
 - using commands in 115
- skipmigrated 110
- skipmigratedoption
 - overview 63
- space management 19
 - adding to a file system 128
 - clusters on GPFS 31
 - files
 - candidates file 187
 - logdir directory 187
 - orphan.stubs file 187
 - premgrdb.dir file 187
 - premgrdb.pag file 187
 - status file 187
 - installing 5
 - maintaining 128

- space management (*continued*)
 - removing 137
 - updating a file system 128
- space management client overview 1
- space management settings 128
 - displaying 139
 - high threshold 128
 - maxcandidates 128
 - premigration setting 128
 - quota 128
 - stub file size 128
- space monitor daemon 83
 - starting 155
- space requirement
 - HSM control files 188
- space usage information
 - displaying 124
- space usage, displaying 114, 123
 - hsmenableimmediatemigrate 95
 - hsmgroupedmigrate 97
 - young files are omitted 95
- spacemgtechnique 25
- standard option format 113
- standard-query restore 65
- starting
 - space management daemons 144
- status file 187
- stopping space management daemons 144
- storage pool 1
 - dsmautomig command 121
- streaming recall 73
 - minimum streaming file size 44
 - overview 3
- stub file size
 - setting 128
- stub files
 - recreating 153
 - recreating deleted 114
 - size 44
- stubsizes 128
- summary of changes V8.1.2 xiii
- synchronizing client and server files 79
- synchronizing client-server files
 - manual reconciliation 161
- syslog information 174

T

- tape recall, optimized
 - list files 76
 - overview 75
 - procedure 78
- tasks 24
- threshold migration 2, 53, 55
 - start manually 57
 - starting manually 121
- thresholds for migration 42
- trace
 - enable client trace 177
 - enable daemon trace 178
 - monitor daemon 181
 - recall daemon 179
 - scout daemon 180
 - watch daemon 179
- trace classes
 - space management client 176
- transfer HSM management of file system 142, 146

- transparent recall 3, 71
- troubleshooting
 - checking daemons 173
 - commands
 - displaying GPFS information 183
 - displaying HSM information 181
 - displaying operating system information 183
 - common problems and solutions 184
 - dump files 174
 - other tools 184
 - process ID file 173
 - syslog information 174
 - trace classes 176
- truncate
 - migrated files 3
- two-way orphan check 161

U

- update 128
- updating space management settings 128
- upgrade
 - multiple servers
 - dsmMultiServerUpgrade.pl 169
- using commands 113

V

- variables 29
- variables, environment 28
- Version 8.1.2 changes xiii

W

- watch daemon 35, 85
- WPARs, adding HSM 39



Product Number: 5725-X04

Printed in USA