



User's Guide



User's Guide

Note

Before using this information and the product it supports, read the information in "Notices" on page 141.

This edition applies to Version 6.1 of IBM Tivoli Storage Manager and to all subsequent releases and modifications until otherwise indicated in new editions or technical newsletters.

© **Copyright International Business Machines Corporation 1995, 2009.**

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

Contents

Figures	vii
--------------------------	------------

Tables	ix
-------------------------	-----------

Preface	xi
--------------------------	-----------

Who should read this publication	xi
Publications	xi
Tivoli Storage Manager publications	xi
Support information	xiii
Getting technical training	xiii
Searching knowledge bases	xiii
Contacting IBM Software Support	xv
Conventions used in this manual	xvi

New for IBM Tivoli Storage Manager Version 6.1.	xix
--	------------

Chapter 1. Space management client for UNIX and Linux overview	1
---	----------

Migrating files overview	2
Premigrating files overview	2
Recalling migrated files overview	3
Reconciling file systems overview	3

Chapter 2. Installing the HSM for UNIX and Linux client	5
--	----------

General installation prerequisites and considerations	5
AIX GPFS HSM client installation overview	7
AIX GPFS HSM client environment requirements	8
Installing the AIX HSM client	8
Upgrading the AIX GPFS HSM client	9
Uninstalling the AIX GPFS HSM client	10
AIX JFS2 HSM client installation overview	10
AIX JFS2 HSM client environment requirements	11
Installing the AIX HSM client	12
Migrating or upgrading the AIX JFS HSM client	12
HP-UX IA64 HSM client installation overview	13
HP-UX IA64 HSM client environment requirements	13
Installing the HP-UX IA64 HSM client	13
Uninstalling the HP-UX IA64 HSM client	13
Linux x86/x86_64 GPFS HSM client installation overview	14
Linux x86/x86_64 GPFS HSM client environment requirements	15
Installing the Linux x86/x86_64 GPFS client	16
Upgrading the Linux x86/x86_64 GPFS HSM client	16
Uninstalling the Linux x86/x86_64 GPFS HSM client	17
Solaris HSM client installation overview	18
Solaris HSM client environment requirements	19
Installing the Solaris HSM client (SPARC only)	19
Solaris local zones support	20

Upgrading the Solaris HSM client	20
Uninstalling the Solaris HSM client	21
Registering your workstation with a Tivoli Storage Manager server	21
Open registration	21
Closed registration	22
Managing your password	22
The HSM GUI	23
Verifying the HSM agent is running	24
Using the HSM GUI	25

Chapter 3. Selecting HSM client options	27
--	-----------

Displaying HSM client options	28
Editing the options files using the backup-archive client	28
Editing dsm.sys options	29
Editing dsm.opt options	32
Optional setup features	32
Assigning management classes to files	33
Displaying management class information	34
Modifying the include-exclude options file	35
Setting environment variables	37
Setting up LAN-free data transfer for HSM.	38
Setting up space management for HACMP on AIX JFS2	39
Space management for AIX GPFS and Linux x86/x86_64 GPFS clusters	42
Configuring the HSM space management agent	44

Chapter 4. Adding and configuring space management	47
---	-----------

Adding space management to file systems	48
Adding space management to nested file systems	49
Mounting the parent file system before nested file system.	49
Adding space management to an exported file system	49
Adding space management to WPARS on AIX 6.1	50
Updating settings from the command line	50
Updating settings from the HSM GUI	50
Space management settings	51
Minimum migration file size	52
Migration threshold percentages	52
Premigration percentage	53
Quotas	54
Stub file size	54
Minimum stream file size.	55
Minimum partial file size.	55
Deactivating space management	55
Reactivating space management	56
Removing space management	57

Chapter 5. Migrating files.	59
--	-----------

Automatic file migration	60
------------------------------------	----

Candidate selection for automatic migration	61
File premigration	62
Manually starting threshold migration	62
Migrating selectively using the dsmmigrate command	63
GPFS policy-driven migration	63

Chapter 6. Archive or retrieve and backup or restore overview 67

Archiving and retrieving files using the backup-archive client	68
Archiving migrated or premigrated files	68
Retrieving archived files	68
Backing up files before or after migration	69
Backing up and restoring premigrated files	69
Restoring files	70
Recreating stub files using the dsmmigundelete command	70
Restoring file systems overview	72
Restoring a file system backed up and migrated to the same server	72
Restoring a file system backed up and migrated to a different server.	73
Restoring a disk	74
Restoring your operating system and file systems	74

Chapter 7. Recalling migrated files. 75

Transparent recall	75
Selective recall	76
Normal recall mode	76
Partial file recall mode.	76
Streaming recall mode.	77
How HSM determines which recall mode to use	78
Setting the recall mode using the dsmmatr command	78

Chapter 8. Reconciling file systems 79

Automatic reconciliation	79
Manually reconciling file systems	80
Resolving orphaned stub files	81

Chapter 9. Space management daemons 83

The space monitor daemon	83
The recall daemon	83
The scout daemon	84
The watch daemon	84
The root daemon	85
Stopping the space management daemons	85

Chapter 10. Scheduling services. 87

Scheduling options	87
Starting the client scheduler	87
Displaying scheduled services information	88

Chapter 11. HSM client dsm.sys and dsm.opt option reference. 89

Candidatesinterval	89
Checkfororphans	90
Checkthresholds.	90

Compression	91
Defaultserver	91
Errorprog	92
Filelist	92
Inclxcl.	93
Maxcandprocs	94
Maxmigrators	94
Maxrecalldaemons	94
Maxreconcileproc	95
Maxthresholdproc	95
Migfileexpiration	96
Migrateserver.	96
Minnmigfilesize	97
Minrecalldaemons	97
Optionformat.	98
Reconcileinterval	98
Restoremigstate	99

Chapter 12. HSM client command reference 101

Standard and short option formats	101
Displaying command-line help	102
Displaying file and file system information	103
HSM command summary	103
dmkilld	105
dsmattr	105
dsmautomig.	107
dsmdf	108
dsmdu	109
dsmls	111
dsmmigs add and update	112
dsmmigs deactivate, reactivate, and remove	116
dsmmigs query	117
dsmmigs sdrreset, enablefailover, disablefailover	118
dsmmigs globaldeactivate, globalreactivate	119
dsmmigs rollback.	120
dsmmigs stop, start, and restart	120
dsmmigs takeover	121
dsmmighelp.	122
dsmmigquery	122
dsmmigrate	124
dsmmigundelete	126
dsmmonitord	127
dsmq	127
dsmrecall.	128
dsmrecalld	130
dsmreconcile	130
dsmrm	131
dsmrootd.	132
dsmscoutd	132
dsmsetpw	133
dsmwatchd	134

Appendix A. The SpaceMan directory 135

Appendix B. HSM .SpaceMan files and subdirectories 137

**Appendix C. Accessibility features for
Tivoli Storage Manager 139**

Notices 141

Trademarks 143

Glossary 145

Index 167

Figures

1. start_HSM : sample script to start HSM	40	3. XML configuration file	44
2. stop_HSM : Stops HSM	41		

Tables

1.	Tivoli Storage Manager server publications	xi	19.	dsmmigundelete command actions without the expiring option	71
2.	Tivoli Storage Manager storage agent publications	xii	20.	dsmmigundelete command actions with the expiring option	71
3.	Tivoli Storage Manager client publications	xii	21.	HSM client determining recall modes	78
4.	Tivoli Storage Manager Data Protection publications	xiii	22.	Automatic reconciliation tasks	79
5.	AIX GPFS HSM client installation packages	8	23.	Option format examples: options without values	102
6.	HSM AIX JFS2 installation packages	11	24.	Option format examples: options with values	102
7.	Linux GPFS available packages	15	25.	HSM commands to display file and file system information	103
8.	Package names and descriptions	16	26.	HSM command summary	103
9.	Linux x86/x86_64 GPFS uninstalling steps	17	27.	HSM management for locally and remotely mounted file systems	110
10.	Solaris package names and directories	18	28.	HSM management for locally and remotely mounted file systems	111
11.	Space Management Options in Your Options Files	30	29.	Space management option settings	114
12.	HSM client options in your dsm.opt file	32	30.	HSM Control Files	135
13.	HSM client attributes in a management class	33	31.	Control files stored in the .SpaceMan directory	137
14.	Assigning management classes to your files	33			
15.	Include and Exclude Statements	35			
16.	Environment Variables	37			
17.	LAN-free data transfer options	38			
18.	HSM agent options	45			

Preface

This publication provides information to help you perform the following tasks:

- Install the HSM client
- Register the HSM client with a Tivoli® Storage Manager server
- Define migration jobs
- Execute and schedule migration jobs
- Search and retrieve migrated and archived files
- Search and restore migrated and backed up files
- Use tools for problem analysis
- Set log levels for problem analysis

Who should read this publication

This publication is intended for those who are responsible for installing, setting up, and administering the HSM client.

Publications

Tivoli Storage Manager publications and other related publications are available online.

You can search all publications in the Tivoli Storage Manager Information Center: <http://publib.boulder.ibm.com/infocenter/tsminfo/v6>.

You can download PDF versions of publications from the Tivoli Storage Manager Information Center or from the IBM® Publications Center at <http://www.ibm.com/shop/publications/order/>.

You can also order some related publications from the IBM Publications Center Web site. The Web site provides information for ordering publications from countries other than the United States. In the United States, you can order publications by calling 800-879-2755.

Tivoli Storage Manager publications

Publications are available for the server, storage agent, client, and Data Protection.

Table 1. Tivoli Storage Manager server publications

Publication title	Order number
<i>IBM Tivoli Storage Manager Messages</i>	GC23-9787
<i>IBM Tivoli Storage Manager Performance Tuning Guide</i>	GC23-9788
<i>IBM Tivoli Storage Manager Problem Determination Guide</i>	GC23-9789
<i>IBM Tivoli Storage Manager for AIX Installation Guide</i>	GC23-9781
<i>IBM Tivoli Storage Manager for AIX Administrator's Guide</i>	SC23-9769
<i>IBM Tivoli Storage Manager for AIX Administrator's Reference</i>	SC23-9775
<i>IBM Tivoli Storage Manager for HP-UX Installation Guide</i>	GC23-9782
<i>IBM Tivoli Storage Manager for HP-UX Administrator's Guide</i>	SC23-9770

Table 1. Tivoli Storage Manager server publications (continued)

Publication title	Order number
<i>IBM Tivoli Storage Manager for HP-UX Administrator's Reference</i>	SC23-9776
<i>IBM Tivoli Storage Manager for Linux Installation Guide</i>	GC23-9783
<i>IBM Tivoli Storage Manager for Linux Administrator's Guide</i>	SC23-9771
<i>IBM Tivoli Storage Manager for Linux Administrator's Reference</i>	SC23-9777
<i>IBM Tivoli Storage Manager for Sun Solaris Installation Guide</i>	GC23-9784
<i>IBM Tivoli Storage Manager for Sun Solaris Administrator's Guide</i>	SC23-9772
<i>IBM Tivoli Storage Manager for Sun Solaris Administrator's Reference</i>	SC23-9778
<i>IBM Tivoli Storage Manager for Windows Installation Guide</i>	GC23-9785
<i>IBM Tivoli Storage Manager for Windows Administrator's Guide</i>	SC23-9773
<i>IBM Tivoli Storage Manager for Windows Administrator's Reference</i>	SC23-9779
<i>IBM Tivoli Storage Manager Server Upgrade Guide</i>	SC23-9554
<i>IBM Tivoli Storage Manager for System Backup and Recovery Installation and User's Guide</i>	SC32-6543

Table 2. Tivoli Storage Manager storage agent publications

Publication title	Order number
<i>IBM Tivoli Storage Manager for SAN for AIX Storage Agent User's Guide</i>	SC23-9797
<i>IBM Tivoli Storage Manager for SAN for HP-UX Storage Agent User's Guide</i>	SC23-9798
<i>IBM Tivoli Storage Manager for SAN for Linux Storage Agent User's Guide</i>	SC23-9799
<i>IBM Tivoli Storage Manager for SAN for Sun Solaris Storage Agent User's Guide</i>	SC23-9800
<i>IBM Tivoli Storage Manager for SAN for Windows Storage Agent User's Guide</i>	SC23-9553

Table 3. Tivoli Storage Manager client publications

Publication title	Order number
<i>IBM Tivoli Storage Manager for UNIX and Linux: Backup-Archive Clients Installation and User's Guide</i>	SC23-9791
<i>IBM Tivoli Storage Manager for Windows: Backup-Archive Clients Installation and User's Guide</i>	SC23-9792
<i>IBM Tivoli Storage Manager for Space Management for UNIX and Linux: User's Guide</i>	SC23-9794
<i>IBM Tivoli Storage Manager for HSM for Windows Administration Guide</i>	SC23-9795
<i>IBM Tivoli Storage Manager Using the Application Program Interface</i>	SC23-9793
<i>Program Directory for IBM Tivoli Storage Manager z/OS Edition Backup-Archive Client</i>	GI11-8912
<i>Program Directory for IBM Tivoli Storage Manager z/OS Edition Application Program Interface</i>	GI11-8911

Table 4. Tivoli Storage Manager Data Protection publications

Publication title	Order number
<i>IBM Tivoli Storage Manager for Advanced Copy Services: Data Protection for Snapshot Devices Installation and User's Guide</i>	SC33-8331
<i>IBM Tivoli Storage Manager for Databases: Data Protection for Microsoft SQL Server Installation and User's Guide</i>	SC32-9059
<i>IBM Tivoli Storage Manager for Databases: Data Protection for Oracle for UNIX and Linux Installation and User's Guide</i>	SC32-9064
<i>IBM Tivoli Storage Manager for Databases: Data Protection for Oracle for Windows Installation and User's Guide</i>	SC32-9065
<i>IBM Tivoli Storage Manager for Enterprise Resource Planning: Data Protection for SAP Installation and User's Guide for DB2</i>	SC33-6341
<i>IBM Tivoli Storage Manager for Enterprise Resource Planning: Data Protection for SAP Installation and User's Guide for Oracle</i>	SC33-6340
<i>IBM Tivoli Storage Manager for Mail: Data Protection for Lotus Domino® for UNIX, Linux, and OS/400® Installation and User's Guide</i>	SC32-9056
<i>IBM Tivoli Storage Manager for Mail: Data Protection for Lotus Domino for Windows Installation and User's Guide</i>	SC32-9057
<i>IBM Tivoli Storage Manager for Mail: Data Protection for Microsoft Exchange Server Installation and User's Guide</i>	SC23-9796
<i>Program Directory for IBM Tivoli Storage Manager for Mail (Data Protection for Lotus Domino)</i>	GI11-8909

Support information

You can find support information for IBM products from a variety of sources.

Getting technical training

Information about Tivoli technical training courses is available online.

Go to <http://www.ibm.com/software/tivoli/education/>.

Searching knowledge bases

If you have a problem with Tivoli Storage Manager, there are several knowledge bases that you can search.

You can begin with the Tivoli Storage Manager Information Center at <http://publib.boulder.ibm.com/infocenter/tsminfo/v6>. From this Web site, you can search all Tivoli Storage Manager publications.

Searching the Internet

If you cannot find an answer to your question in the Tivoli Storage Manager information center, search the Internet for the latest, most complete information that might help you resolve your problem.

To search multiple Internet resources, go to the support Web site for Tivoli Storage Manager at <http://www.ibm.com/software/sysmgmt/products/support/IBMTivoliStorageManager.html>. From there, you can search a variety of resources including:

- IBM technotes
- IBM downloads

- IBM Redbooks®

If you still cannot find the solution to the problem, you can search forums and newsgroups on the Internet for the latest information that might help you resolve your problem. To share your experiences and learn from others in the user community, go to the Tivoli Storage Manager wiki at <http://www.ibm.com/developerworks/wikis/display/tivolistoragemanager/Home>.

Using IBM Support Assistant

At no additional cost, you can install on any workstation the IBM Support Assistant, a stand-alone application. You can then enhance the application by installing product-specific plug-in modules for the IBM products that you use.

The IBM Support Assistant helps you gather support information when you need to open a problem management record (PMR), which you can then use to track the problem. The product-specific plug-in modules provide you with the following resources:

- Support links
- Education links
- Ability to submit problem management reports

For more information, see the IBM Support Assistant Web site at <http://www.ibm.com/software/support/isa/>.

Finding product fixes

A product fix to resolve your problem might be available from the IBM Software Support Web site.

You can determine what fixes are available by checking the Web site:

1. Go to the IBM Software Support Web site at <http://www.ibm.com/software/tivoli/products/storage-mgr/product-links.html>.
2. Click the **Support Pages** link for your Tivoli Storage Manager product.
3. Click **Download**, and then click **Fixes by version**.

Getting e-mail notification of product fixes

You can get notifications about fixes and other news about IBM products.

To receive weekly e-mail notifications about fixes and other news about IBM products, follow these steps:

1. From the support page for any IBM product, click **My support** in the upper-right corner of the page.
2. If you have already registered, skip to the next step. If you have not registered, click **Register** in the upper-right corner of the support page to establish your user ID and password.
3. Sign in to **My support**.
4. On the My support page, click **Edit profiles** in the left navigation pane, and scroll to **Select Mail Preferences**. Select a product family and check the appropriate boxes for the type of information you want.
5. Click **Submit**.
6. For e-mail notification for other products, repeat steps 4 and 5.

Contacting IBM Software Support

You can contact IBM Software Support if you have an active IBM software maintenance contract and if you are authorized to submit problems to IBM.

Before you contact IBM Software Support, follow these steps:

1. Set up a software maintenance contract.
2. Determine the business impact of your problem.
3. Describe your problem and gather background information.

Then see “Submit the problem to IBM Software Support” on page xvi for information on contacting IBM Software Support.

Setting up a software maintenance contract

Set up a software maintenance contract. The type of contract that you need depends on the type of product you have.

- For IBM distributed software products (including, but not limited to, Tivoli, Lotus[®], and Rational[®] products, as well as IBM DB2[®] and IBM WebSphere[®] products that run on Microsoft[®] Windows[®] or UNIX[®] operating systems), enroll in IBM Passport Advantage[®] in one of the following ways:
 - **Online:** Go to the Passport Advantage Web page at <http://www.ibm.com/software/lotus/passportadvantage/>, click **How to enroll**, and follow the instructions.
 - **By Phone:** For the phone number to call in your country, go to the IBM Software Support Handbook Web page at <http://www14.software.ibm.com/webapp/set2/sas/f/handbook/home.html> and click **Contacts**.
- For server software products, you can purchase a software maintenance agreement by working directly with an IBM sales representative or an IBM Business Partner. For more information about support for server software products, go to the IBM Technical support advantage Web page at <http://www.ibm.com/servers/>.

If you are not sure what type of software maintenance contract you need, call 1-800-IBMSERV (1-800-426-7378) in the United States. For a list of telephone numbers of people who provide support for your location, go to the Software Support Handbook page at <http://www14.software.ibm.com/webapp/set2/sas/f/handbook/home.html>.

Determine the business impact

When you report a problem to IBM, you are asked to supply a severity level. Therefore, you need to understand and assess the business impact of the problem you are reporting.

Severity 1	Critical business impact: You are unable to use the program, resulting in a critical impact on operations. This condition requires an immediate solution.
Severity 2	Significant business impact: The program is usable but is severely limited.
Severity 3	Some business impact: The program is usable with less significant features (not critical to operations) unavailable.
Severity 4	Minimal business impact: The problem causes little impact on operations, or a reasonable circumvention to the problem has been implemented.

Describe the problem and gather background information

When explaining a problem to IBM, it is helpful to be as specific as possible. Include all relevant background information so that IBM Software Support specialists can help you solve the problem efficiently.

To save time, know the answers to these questions:

- What software versions were you running when the problem occurred?
- Do you have logs, traces, and messages that are related to the problem symptoms? IBM Software Support is likely to ask for this information.
- Can the problem be recreated? If so, what steps led to the failure?
- Have any changes been made to the system? For example, hardware, operating system, networking software, and so on.
- Are you currently using a workaround for this problem? If so, be prepared to explain it when you report the problem.

Submit the problem to IBM Software Support

You can submit the problem to IBM Software Support online or by phone.

Online

Go to the IBM Software Support Web site at <http://www.ibm.com/software/support/probsub.html>. Enter your information into the appropriate problem submission tool.

By phone

For the phone number to call in your country, go to the contacts page of the IBM Software Support Handbook at <http://www14.software.ibm.com/webapp/set2/sas/f/handbook/home.html>.

If the problem that you submit is for a software defect or for missing or inaccurate documentation, IBM Software Support creates an Authorized Program Analysis Report (APAR). The APAR describes the problem in detail. If a workaround is possible, IBM Software Support provides one for you to implement until the APAR is resolved and a fix is delivered. IBM publishes resolved APARs on the Tivoli Storage Manager product support Web site at <http://www.ibm.com/software/sysmgmt/products/support/IBMTivoliStorageManager.html>, so that users who experience the same problem can benefit from the same resolutions.

Conventions used in this manual

This manual 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 variables or 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.

timeformat Boldface italic type indicates a Tivoli Storage Manager option. The bold type is used to introduce the option, or used in an example.

Occasionally, file names are entered in boldface italic for emphasis.

Example	Description
<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.
<code>maxcmdretries</code>	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 Tivoli Storage Manager Version 6.1

The following features are new for IBM Tivoli Storage Manager in Version 6.1:

New in Version 6.1.3

- **HSM Configuration Enhancements**

- The HSM configuration enhancements enable the administrator to reduce the space consumed by the HSM configurable complete file index (CFI) size. These enhancements also help the administrator avoid setting unrealistic low migration thresholds.

Refer to the `dsmmigfs` command **MAXFiles** parameter for details.

New in Version 6.1

- **GPFS™ 3.2 storage pool support**

- Tivoli Storage Manager 6.1 allows multiple General Parallel File System (GPFS) storage pools in one file system. Monitoring a file system includes monitoring each storage pool in the file system.
- The automigration command, **dmsautomig**, permits automatic migration of storage pools and file systems.

- **Hierarchical Storage Management (HSM) for AIX® and Linux® for GPFS does not require RSCT file set for cluster support**

A new responsiveness service function is being added. It provides:

- node response monitoring
- node failure detection and initiate failover actions
- event notification processing

This function replaces the requirement of installing the RSCT Group Services. Do not use RSCT file set for cluster support.

- **Partial file recall enhancements for optimal tape access**

The **dsmrecall** command recalls partial files with the *-OFFset* and *-SIZE* options. With this command, you can specify the portion of a file that is to be recalled.

Chapter 1. Space management client for UNIX and Linux overview

The IBM Tivoli Storage Manager for Space Management client for UNIX and Linux (the HSM client) migrates files from your local file system to distributed 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 HSM client functions for threshold migration, demand migration, selective migration, selective and transparent recall now includes processing GPFS file systems containing multiple HSM managed storage pools.

The HSM client has both a graphical user interface (the HSM GUI) and commands 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 owned by user ibm:

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

Your Tivoli Storage Manager administrator associates management classes with your files. You, as root user:

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

The options and settings you define for the HSM client determine which files are eligible for automatic migration, the order in which files are migrated, where the migrated files are stored, and 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.

When a file is migrated from your local system to Tivoli Storage Manager 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 so that the files appear to reside locally. This contrasts with archiving, where you usually delete files from your local file system after archiving them.

The HSM 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 reside on your local file system or are migrated to Tivoli Storage Manager storage. You can use the IBM Tivoli Storage Manager backup-archive client to back up and restore migrated files in the same manner as you would back up and restore files that reside on your local file system. 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.

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. You can use the backup-archive client to archive and retrieve copies of migrated files in the same manner as you would archive and retrieve copies of files that reside on your local file system.

Migrating files overview

The HSM client provides both automatic and selective migration. Once file migration begins, the HSM client sends a copy of your file to 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 HSM 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 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. For HSM on AIX JFS2, HP-UX and Solaris, as files are migrated, space becomes available on your file system and the process or event that caused the out-of-space condition continues. On AIX GPFS and Linux x86/x86_64 GPFS, the process receives an out-of-space (ENOSPC) return code and stops.

Selective migration moves specific files from your local file system to 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 storage to free additional space on your local file system.

Premigrating files overview

For faster migration, you can use the premigration process to prepare your files for automatic migration. Premigrated files are copied to storage while the original files remain on your local file system. Then, the next time you need free space on your local file system, the HSM client just changes the status of premigrated files to migrated files without requiring additional time to copy them to storage.

As part of the process, the HSM client verifies that files have not changed since they were premigrated and replaces the local file system copies of the files with stub files.

Files are premigrated each time automatic migration completes if:

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

Recalling migrated files overview

You can recall a migrated file to your local file system from 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. See “Selective recall” on page 76.

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 HSM 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 storage remains valid. The local copy is premigrated in this case.

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

Partial file recall mode recalls a portion of a migrated file and is valid for AIX GPFS and Linux x86/x86_64 GPFS only. This avoids having to recall an entire, potentially large file, when only a small portion of the file is required by an application. When HSM intercepts a read request for a file configured for partial file recall, it will calculate 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. See “Partial file recall mode” on page 76.

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. See “Streaming recall mode” on page 77.

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

Reconciling file systems overview

When you modify the data of a migrated or premigrated file or erase a migrated or premigrated file from your local file system, you retain an obsolete copy of the file in storage. 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. You, as root user, can change this by setting the *reconcileinterval* option to a value larger than 0. 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.

Chapter 2. Installing the HSM for UNIX and Linux client

You install the HSM client on your workstation and register it as a client node with a Tivoli Storage Manager server. The files on your node must be associated with a management class on the Tivoli Storage Manager server which is configured for space management. Make sure you read both the general and system-specific requirements before installing the HSM client.

You can add space management to the following file systems:

- General Parallel File System (GPFS) on AIX cluster
- General Parallel File System (GPFS) on Linux x86/x86_64 cluster
- Enhanced Journaled File System (JFS2) on an AIX workstation
- Veritas File System (VxFS) on a Sun Solaris workstation
- OnlineJFS or Veritas File System (VxFS) on a HP-UX-IA64

General installation prerequisites and considerations

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

Installation prerequisites

- You must have root user authority to install, set up, and use the HSM client on your workstation.
- You must install the API *before* you install your operating system's backup-archive client, and you must install and set up the backup-archive client *before* you install the HSM client. Both the HSM and backup-archive clients share common code, the same options files, communication protocols, node registration, and storage.

Note: Refer to the Backup-Archive Clients Installation and User's Guide for your platform for supported languages and locales.

Warnings and recommendations

The `/usr/tivoli/tsm/client/hsm/bin` directory is one of the base directories into which the HSM 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

1. Prior to installing anything, read *all* of these general requirements and then read your system's installation overview, specific environment requirements, and each installation step to ensure you are ready to install the HSM client.
2. Follow your system specific installation procedure.

3. When you complete the installation, register your workstation as a node with a TMS server.
4. Modify the `dsm.sys` and `dsm.opt` configuration files (see Chapter 3, “Selecting HSM client options,” on page 27). Both files are located in the `/usr/tivoli/tsm/client/ba/bin` directory or the `/opt/tivoli/tsm/client/ba/bin` directory for HP-UX, Linux, and Solaris.

Reinstallation or upgrade overview

Before you *reinstall* the HSM client, stop all activity and do not access any files on file systems to which you added space management. The install process will fail 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 in order to get those packages at the same level as the HSM client.

HSM GUI requirements and considerations

- The space management agent and the HSM GUI versions must match, otherwise the authentication with the space management agent is not possible.
- In order to use the HSM GUI, the Tivoli Storage Manager node name and password are required. The node name and password are the same the HSM client uses to authenticate with the Tivoli Storage Manager server.
- Before you can connect to a Tivoli Storage Manager client node using the HSM GUI, that client node must be registered and authenticated with a Tivoli Storage Manager server.
- The HSM GUI requires the following software in order to run:
 - The API, backup-archive client, and HSM client
 - The Java™ Runtime Environment (JRE) 1.5.x or JRE 1.6.x where x is 1 or later.
- You can access Internet resources by clicking the **Help** menu, by clicking the top banner area or by downloading the required version of the JRE.
- The PATH environment variable must include the path to the java executable. In order to verify you have the right JRE installed, issue the **java -version** command from a UNIX terminal window from the system on which you want to start the HSM GUI.
- If a lower version of JRE is detected, a window opens with warning information and a button to open a browser to the location where the JRE can be downloaded. The link provided might be different, depending on the platform where the HSM GUI is started. However, you can manually download the required version of the JRE from the following URLs:
 - JRE for AIX platforms: <http://www-128.ibm.com/developerworks/java/jdk/aix/service.html>
 - JRE for HP platforms: <http://www.hp.com/products1/unix/java/index.html>
 - J2SE 5.0 or 6.0 for Linux, Solaris, and Windows platforms: <http://java.sun.com/j2se>

AIX GPFS HSM client installation overview

Before installing the HSM client on AIX GPFS file 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:

- The HSM AIX GPFS client is not compatible with the HSM AIX JFS2 client or the backup-archive client for JFS2. If you have either of these clients installed and want to install the HSM AIX GPFS client, you need to remove the JFS2 clients.
- The HSM client for AIX GPFS supports GPFS 3.2 cluster with GPFS file systems.
- The HSM client only manages file systems belonging to the local (home) GPFS cluster; it does not manage remotely mounted file systems.
- On AIX 6.1, the HSM client can be installed in the global partition and supports transparent recall for both global and local WPARs. Using HSM commands from a local WPAR is not supported. For more information on WPARs see “Adding space management to WPARS on AIX 6.1” on page 50.
- In a GPFS environment, a small file that is less than a GPFS block size can become larger after an HSM migration because GPFS adds meta information to the file during the migration. Because another block on the file system is allocated for the meta information, this increases the space allocated for the file. If a file system is filled to its maximum capacity with many small files, it is possible that the file system can run out of space during the file migration.

When you install the HSM client on GPFS file systems, the install process:

- 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** command 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.

AIX GPFS installation overview

For an initial installation, follow these steps exactly in this order:

1. Edit the `dsm.opt` and `dsm.sys` files installed with the backup-archive client to define the basics for using Tivoli Storage Manager HSM (see “Editing the options files using the backup-archive client” on page 28).
2. Install the HSM client on each HSM node (see “Installing the AIX HSM client” on page 8).
3. Make sure that after installation, the `dsmrecalld` daemon is running on at least one node.
4. Unmount all GPFS file systems on all nodes within the GPFS cluster that will be HSM managed.

5. Issue the following command to activate DMAPI management for these GPFS file systems:

```
mmchfs -z yes <device>
```
6. Remount all GPFS file systems on all nodes within the GPFS cluster. The HSM daemons will detect each node's initial state and assign all nodes an instance number in relation to the GPFS cluster definition.
7. On the HSM owner nodes, add HSM management to all GPFS file systems by issuing the command.

```
dsmmigfs add <fs>
```
8. Enable failover of HSM on all cluster nodes that participate in distributed HSM (owner and source nodes) with the **dsmmigfs enablefailover** command.

```
dsmmigfs enablefailover
```

Table 5 displays the packages available on the installation media in the `/usr/sys/inst.images` directory:

Table 5. AIX GPFS HSM client installation packages

Package	Installs	Into
tivoli.tsm.client.ba.gpfs	The backup-archive client for AIX GPFS	The <code>/usr/tivoli/tsm/client/ba/bin</code> directory
tivoli.tsm.client.hsm.gpfs	The HSM client for AIX GPFS	The <code>/usr/tivoli/tsm/client/hsm/bin</code> directory
tivoli.tsm.client.api	The API for AIX	The <code>/usr/tivoli/tsm/client/api/bin</code> directory

AIX GPFS HSM client environment requirements

Before you install the HSM client, your workstation must meet minimum hardware and software requirements.

Hardware requirements

The AIX GPFS HSM client requires the RS/6000® (32-bit or 64-bit) pSeries® or compatible hardware.

Software requirements

The AIX GPFS HSM client requires the following software in order to run:

- AIX 5L™ 5.3 ML 1 PPC in 32 or 64-bit kernel mode or AIX 5L 6.1 PPC in 32 or 64-bit kernel mode
- GPFS 3.2

Installing the AIX HSM client

You can install the AIX HSM client directly from the CD-ROM or from a local directory where you copy the client files.

Follow these steps to install the HSM client from the AIX command line or CD-ROM:

Note: If you are installing from a hard disk directory, remove the `.toc` file from that directory before issuing the **smitty** command.

1. Issue the **smitty install** command.

2. Select **Install and Update Software** and press **Enter**.
3. Select **Install and Update from ALL Available Software** and press **Enter**.
4. In the **INPUT device / directory for software** field, select one of the following and press **Enter**:
 - a. `/usr/sys/inst.images` (directory that installation images reside)
 - b. `/dev/cd0` (for a CD-ROM install)
5. In the **SOFTWARE to Install** field press **F4**.
6. Scroll through the file sets and press **F7** to select the Tivoli Storage Manager file sets that you want to install and press **Enter**. The HSM client installation requires a minimum of:
 - Tivoli Storage Manager application programming interface
 - Tivoli Storage Manager backup-archive client common files
 - Tivoli Storage Manager backup-archive client base files
7. Select the options that you want and press **Enter** to begin the installation.
8. If needed, unmount your CD-ROM drive.
9. Continue with “Registering your workstation with a Tivoli Storage Manager server” on page 21.

Upgrading the AIX GPFS HSM client

To upgrade the HSM client, and then upgrade from GPFS Version 2.2 or lower, you need to perform a series of steps.

Before upgrading the HSM client, we recommend you read the AIX GPFS installation overview (see “AIX GPFS HSM client installation overview” on page 7).

If you are upgrading GPFS from version 2.2 or lower:

- Depending on the cluster type you have used so far, you might need to configure the cluster file system.
- Follow the GPFS migration procedure in *GPFS: Concepts, Planning, and Installation Guide, Version 3.2*.

If you are upgrading GPFS to version 3.2, upgrade all cluster nodes to version 3.2 and set the `dmapifilehandlesize=32`. If these tasks are not done, the HSM client might fail with the following errors:

- The `dsmautomig` processes do not start, and this message is issued:
no candidates found in file system <file-system>
- The `APool` file in the `/etc/adsm/SpaceMan/candidatesPool/APool.filesystem` directory is created with a size of 0.

To check if the `dmapifilehandlesize` parameter is set to the correct value, issue the following command:

```
mmfsadm dump config | grep -i dmapifilehandlesize
```

If the output from the command is `dmapifilehandlesize 16`, set the `dmapifilehandlesize=32`.

To upgrade the HSM client, perform the following steps in order:

1. Ensure that all HSM nodes (owner and backup) are in a consistent state and that all HSM 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 Tivoli Storage Manager for Space Management from all nodes.
6. Install Tivoli Storage Manager for Space Management on all nodes as described in “AIX GPFS HSM client installation overview” on page 7 for the initial installation.
7. Globally reactivate HSM on every node by issuing the command: **dsmmigfs globalreactivate**.
8. Enable failover on every node with 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 will also run on these nodes, which are designated for failover feature.

Uninstalling the AIX GPFS HSM client

Follow this procedure to uninstall the AIX GPFS HSM client.

Before uninstalling the HSM client, we recommend you read the AIX GPFS installation overview (see “AIX GPFS HSM client installation overview” on page 7).

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

Follow these steps to uninstall the AIX HSM client:

1. To remove HSM support from all managed file systems, issue the **dsmmigfs remove** command.
2. Issue the **smitty remove** command.
3. In the **SOFTWARE name** field, press **F4** to list the Tivoli Storage Manager file sets that you want to uninstall.
4. Select the Tivoli Storage Manager file sets that you want to uninstall and press **Enter**.
5. In the **PREVIEW only?** (remove operation will *not* occur) field, select **No** and press **Enter**.

AIX JFS2 HSM client installation overview

Before installing the HSM client on AIX JFS2, you need to review the general requirements, system specific requirements, and review all installation steps.

Note:

- The HSM AIX JFS2 client is not compatible with the HSM AIX GPFS client or the backup-archive client for GPFS. If you have either of these clients installed and want to install the HSM AIX JFS2 client, you need to remove the GPFS clients.
- The procedure to upgrade an AIX JFS HSM client is described in “Migrating or upgrading the AIX JFS HSM client” on page 12.

- The HSM client is not supported on JFS2 Encrypted File Systems (EFS). For more information see “dsmmigfs add and update” on page 112.
- Extended attributes (EA_{v2}) are supported by the HSM client. However, if you delete a stub file after that file was migrated and use the **dsmmigundelete** command to recreate the stub file, the extended attributes are not restored.
- On AIX 6.1, the HSM client can be installed in the global partition and supports transparent recall for both global and local WPARs. Using HSM commands from a local WPAR is not supported. For more information on WPARs see “Adding space management to WPARs on AIX 6.1” on page 50.

Table 6 lists the packages that are available on the installation media:

Table 6. HSM AIX JFS2 installation packages

Package	Installs	Into the following directory
tivoli.tsm.client.ba	The backup-archive client for AIX	/usr/tivoli/tsm/client/ba/bin
tivoli.tsm.client.hsm	The HSM client for AIX JFS2 64-bit	/usr/tivoli/tsm/client/hsm/bin
tivoli.tsm.client.api	The API for AIX	/usr/tivoli/tsm/client/api/bin
tivoli.tsm.client.msg.lang	The NL messages for Hierarchical Storage Management. American English messages are already included in the Hierarchical Storage Management code.	/usr/tivoli/tsm/client/lang/bin

AIX JFS2 HSM client environment requirements

Before installing the HSM client on an AIX JFS2 system, you need to review the hardware and software requirements.

Hardware requirements

- RS/6000 (32-bit or 64-bit), pSeries, or compatible hardware
- The HACMP™ support of the client requires the following:
 - At least two identical pSeries (RS/6000)
 - Two or more network adapters for each machine

Software requirements

The HSM AIX JFS2 client requires the following software in order to run:

- AIX 5L 5.3 ML 1 PPC in 32 or 64-bit kernel mode or AIX 5L 6.1 PPC in 32 or 64-bit kernel mode

Installing the AIX HSM client

You can install the AIX HSM client directly from the CD-ROM or from a local directory where you copy the client files.

Follow these steps to install the HSM client from the AIX command line or CD-ROM:

Note: If you are installing from a hard disk directory, remove the .toc file from that directory before issuing the **smitty** command.

1. Issue the **smitty install** command.
2. Select **Install and Update Software** and press **Enter**.
3. Select **Install and Update from ALL Available Software** and press **Enter**.
4. In the **INPUT device / directory for software** field, select one of the following and press **Enter**:
 - a. `/usr/sys/inst.images` (directory that installation images reside)
 - b. `/dev/cd0` (for a CD-ROM install)
5. In the **SOFTWARE to Install** field press **F4**.
6. Scroll through the file sets and press **F7** to select the Tivoli Storage Manager file sets that you want to install and press **Enter**. The HSM client installation requires a minimum of:
 - Tivoli Storage Manager application programming interface
 - Tivoli Storage Manager backup-archive client common files
 - Tivoli Storage Manager backup-archive client base files
7. Select the options that you want and press **Enter** to begin the installation.
8. If needed, unmount your CD-ROM drive.
9. Continue with "Registering your workstation with a Tivoli Storage Manager server" on page 21.

Migrating or upgrading the AIX JFS HSM client

If you want to migrate or upgrade from another AIX HSM managed file system type, such as JFS, you must follow a specific procedure.

Note: The **mount** parameter is not required for any options with JFS2. You only need to issue a **dsmmigfs add <fs>** command after the first mount.

Follow these steps:

1. Back up the JFS HSM-managed file systems.
2. Unmount the JFS HSM-managed file systems.
3. Uninstall the JFS HSM client.
4. Mount the JFS2 file system(s) at the same mount point(s) where the JFS HSM-managed JFS file system(s) were mounted.
5. Install the JFS2 HSM client (see "Installing the AIX HSM client" on page 8).
6. Perform a backup-archive client restore operation with the **restoremigstate** option set to *yes*.

HP-UX IA64 HSM client installation overview

Before installing the HSM client on HP-UX IA64, you need to review the environment requirements and review all steps prior to beginning the installation.

HP-UX IA64 HSM client environment requirements

Before you install the HSM client, your workstation must meet minimum hardware and software requirements.

Hardware requirements

- An HP-UX Itanium[®] 2 processor
- Disk space: 550 MB hard disk space
- Memory: 512 MB
- A minimum screen resolution of 800 by 600 pixels is required to display the Tivoli Storage Manager HSM GUI

Software requirements

- HP-UX 11i V2 for Itanium 2

Installing the HP-UX IA64 HSM client

Use a specific command to install the HSM client on HP-UX IA64.

To install the HSM packages issue the following commands:

1. Issue the command:

```
/usr/sbin/swinstall -x mount_all_filesystems=false -v -s `pwd`/TIVsmC TIVsm64
```

2. Issue the command:

```
/usr/sbin/swinstall -x mount_all_filesystems=false -v -s `pwd`/TIVsmChsm TIVsm64
```

Uninstalling the HP-UX IA64 HSM client

If you want to uninstall the HP-UX IA64 HSM client, you must follow a specific set of steps. If you want to upgrade the HSM client, you must uninstall the backup-archive client and the API and then reinstall them at the same version and level as the HSM client. A mixed installation is *not* supported.

Note:

Activity on HSM-managed file systems should be avoided during software installation. Access to migrated files during the upgrade will not be possible, because the system would be suspended while trying to read a file stub.

If you want to completely remove the HSM client, without reinstalling a new version complete all of the following steps. To uninstall the HSM client so you can upgrade the client, skip step 1 and then follow all other steps and then see "Installing the HP-UX IA64 HSM client."

1. To remove HSM support from all managed file systems, issue the command:
dsmmigfs remove.

Note: 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. To remove the HP-UX IA64 HSM client package, issue the following command as root user:

```
/usr/sbin/swremove -x mount_all_filesystems=false -v \ TIVsm64.CLIENT_HSM
```

3. Issue the following command to uninstall all the components of the backup-archive client (command line, GUI, Web GUI, and the administrative client). You cannot uninstall a single component of this package (for example, the Web GUI) without uninstalling the complete package.


```
/usr/sbin/swremove -x mount_all_filesystems=false -v TIVsm64.CLIENT
```
4. If one or more Tivoli Storage Manager language messages packages have been installed, these must be removed before removing the API package. Issue the following command as root user, where `xx_XX` is the language you wish to remove:


```
/usr/sbin/  
swremove -x mount_all_filesystems=false -v \ TIVsm64.CLIENT_msg_xx_XX
```
5. If you also want to remove the API file set, issue the following command:


```
/usr/sbin/swremove -x mount_all_filesystems=false -v \ TIVsm64.CLIENT_API64
```

Linux x86/x86_64 GPFS HSM client installation overview

Before installing the HSM client on Linux x86/x86_64 GPFS file systems, you need to 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:

- The HSM client only manages file systems belonging to the local (home) GPFS cluster; it does not manage remotely mounted file systems.
- HSM cluster installations are certified on IBM Linux Cluster 1350™ (see the IBM Redbooks: *Linux Clustering with CSM and GPFS* SG24-6601).
- Also see the recommendations provided with the IBM GPFS for Linux on x86/x86_64 architecture.
- In a GPFS environment, a small file that is less than 8 KB (smaller than a GPFS block size) can become larger after an HSM migration because GPFS adds meta information to the file during the migration. Because another block on the file system is allocated for the meta information, this increases the space allocated for the file. If a file system is filled to its maximum capacity with many small files, it is possible that the file system can run out of space during the file migration.

When you install the HSM client on Linux x86/x86_64 GPFS file systems, the install process:

- Stops any space management daemons that are running.
- Removes any statement from the `/etc/inittab` file that loads the `dsmswatchd` 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 `dsmswatchd` command 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.
- Creates two HSM-specific files (`DSMNodeset` file and `DSMSDRVersion` file) for each GPFS node set, stores them in the GPFS internal repository, and starts the space management daemons.

Table 7 lists the packages available on the installation media:

Table 7. Linux GPFS available packages

Package	Installs	Into
TIVsm-API.i386.rpm	The API for Linux x86/x86_64	The /opt/tivoli/tsm/client/api/bin directory
TIVsm-BA.i386.rpm	The Tivoli Storage Manager backup-archive client (command line), the administrative client (command line), and the Web backup-archive client	The /opt/tivoli/tsm/client/ba/bin directory
TIVsm-HSM.i386.rpm	The HSM client for Linux x86/x86_64	The /opt/tivoli/tsm/client/hsm/bin directory
TIVsm-msg.lang.i386.rpm	The message package for language (where <i>lang</i> is the language code)	/opt/tivoli/tsm/client/lang

Initial installation overview

1. Edit the dsm.opt and dsm.sys files installed with the backup-archive client to define the basics for using Tivoli Storage Manager HSM.
2. Install Tivoli Storage Manager for Space Management on each HSM node.
3. Make sure that after installation, the dsmrecalld daemon is running on at least one node.
4. For every node that will be HSM managed, unmount all GPFS file systems on all nodes within the GPFS nodeset.
5. Enable the Data Management API (DMAPI) for GPFS for all file systems that will be managed by HSM (this needs to be done only once for each file system):
 - a. Unmount all GPFS file systems on all nodes, within the GPFS cluster, that will be managed by HSM
 - b. Activate DMAPI management for these GPFS file systems:
`mmchfs <device> -z yes`
 - c. Remount all GPFS file systems on all nodes within the GPFS cluster

Refer to *GPFS: Administration and Programming Reference* for further information on GPFS commands.

Linux x86/x86_64 GPFS HSM client environment requirements

Before you install the HSM client on a Linux x86/x86_64 GPFS system, your workstation must meet minimum communication, hardware, software requirements.

Software requirements

The IBM Tivoli Storage Manager HSM client requires the following software in order to run with the following distributions:

- GPFS 3.2 on 32bit SUSE Linux Enterprise Server 10 or Red Hat Enterprise Linux 5.0

Note: X Windows System X11R6 is a requirement to use the HSM GUI. If you do not plan to use the HSM GUI, add the `--nodeps` option of the `rpm` command to disable the check for this requirement.

Installing the Linux x86/x86_64 GPFS client

You can install the HSM client on Linux x86/x86_64 GPFS systems using the product CD-ROM.

Follow these steps:

1. Log in as the root user and mount the CD-ROM to the /cdrom directory.
2. Changed directories to the installation packages directory on the CD:
cdrom/tsmcli/linux86
3. Install the Tivoli Storage Manager clients in the order that is presented in Table 8. During installation, these packages are installed in unique directories. Press **y**) to each question when prompted.
4. Restart your system when you complete the installation.

Table 8. Package names and descriptions

To install	Issue this command
API	To install the API files in the /opt/tivoli/tsm/client/api/bin directory: <pre>rpm -i TIVsm-API.i386.rpm</pre>
Backup-archive client (CLI and GUI) Web client Administrative client	To install the backup-archive client (command line, GUI, and Web client) in the /opt/tivoli/tsm/client/ba/bin directory: <pre>rpm -i TIVsm-BA.i386.rpm</pre> To circumvent the dependence check, you can use the --nodeps option, but then you must check the dependencies manually. To install the backup-archive client (command line, GUI, and Web client) in the /opt/tivoli/tsm/client/ba/bin directory: <pre>rpm -i --nodeps TIVsm-API.i386.rpm</pre> To install the administrative client (command line) in the /opt/tivoli/tsm/client/admin/bin directory: <pre>rpm -i --nodeps TIVsm-BA.i386.rpm</pre>
HSM client	To install the following HSM client files in the /opt/tivoli/tsm/client/hsm/bin directory. <ul style="list-style-type: none">• Tivoli Storage Manager for Space Management commands• Space management agent (hsmagent)• IBM Tivoli Enterprise Space Management Console (HSM GUI) <pre>rpm -i TIVsm-HSM.i386.rpm</pre>

Upgrading the Linux x86/x86_64 GPFS HSM client

To upgrade the HSM client, and then upgrade from GPFS Version 2.2 or lower, you need to perform a series of steps.

Before upgrading the HSM client, we recommend you read the Linux GPFS installation overview (see “Linux x86/x86_64 GPFS HSM client installation overview” on page 14).

If you are upgrading GPFS from version 2.2 or lower:

- Also, follow the GPFS migration procedure from *GPFS: Concepts, Planning, and Installation Guide, Version 3.2*.

If you are upgrading GPFS to version 3.2, upgrade all cluster nodes to version 3.2 and set the `dmapifilehandlesize=32`. If these tasks are not done, the HSM client might fail with the following errors:

- The `dsmautomig` processes do not start, and this message is issued:
no candidates found in file system <file-system>
- The `Apool` file in the `/etc/adsm/SpaceMan/candidatesPool/Apool.filesystem` directory is created with a size of 0.

To check if the `dmapifilehandlesize` parameter is set to the correct value, issue the following command:

```
mmfsadm dump config | grep -i dmapifilehandlesize
```

If the output from the command is `dmapifilehandlesize 16`, set the `dmapifilehandlesize=32`.

To upgrade the HSM client, perform the following steps in order:

1. Ensure that all HSM nodes (owner and backup) are in a consistent state and that all HSM 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 Tivoli Storage Manager for Space Management from all nodes.
6. Install Tivoli Storage Manager for Space Management on all nodes as described in “Linux x86/x86_64 GPFS HSM client installation overview” on page 14 for the initial installation.
7. Globally reactivate HSM on every node by issuing the command: **`dsmmigfs globalreactivate`**.
8. Enable failover on every node with 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 will also run on these nodes, which are designated for failover feature.

Uninstalling the Linux x86/x86_64 GPFS HSM client

You must follow a specific procedure to uninstall the Linux x86/x86_64 GPFS HSM client.

If you want to upgrade the HSM 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 HSM-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.

Use Table 9 to know which steps to run to uninstall the HSM client.

Table 9. Linux x86/x86_64 GPFS uninstalling steps

To accomplish:	Perform steps:
To completely remove the HSM client	1 and 2

Table 9. Linux x86/x86_64 GPFS uninstalling steps (continued)

To accomplish:	Perform steps:
To upgrade the HSM client	2 and 3

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**
3. To install the latest version, refer to “Linux x86/x86_64 GPFS HSM client installation overview” on page 14.

Solaris HSM client installation overview

Before installing the HSM client on Solaris, you need to review the general requirements, system specific requirements, and all installation steps.

Review these considerations before installing the HSM client:

- VxFS 4.1 or 5.0 is required when installing HSM.
- You must remove previous versions of Tivoli Storage Manager or ADSM before installing a new version of the HSM client. Refer to “Uninstalling the Solaris HSM client” on page 21 for instructions on removing previously installed Tivoli Storage Manager packages.
- The Solaris file system requires that the kernel extension (adsmhsm) is loaded before you can use the HSM client. This kernel extension loads automatically during startup. Once the kernel extension is loaded, you cannot unload it. If you plan to install a new version of the HSM client, remove the HSM client and restart your system before you install a new version. If you do not plan to reinstall the HSM client, remove the HSM client and restart your system to unload the kernel extension.
- A Tivoli Storage Manager installation administration file (tsmadmin) is used in place of the default administration file /var/sadm/install/admin), so that you are not asked about setuid, setgid, or superuser permission during installation. If you want to use the default administration file, remove the *-a /tsmadmin* option from the commands listed, and answer the questions about setuid, setgid, or superuser permission during installation with *y*.

Table 10 lists the packages and their installed directory locations for Solaris.

Table 10. Solaris package names and directories

This software package	Is in this directory
TIVsmCapi.pkg	/opt/tivoli/tsm/client/api
TIVsmCba.pkg	/opt/tivoli/tsm/client/ba
TIVsmChsm.pkg	/opt/tivoli/tsm/client/hsm
TIVsmCl<Xx>.pkg	/opt/tivoli/tsm/client/lang/<Xx>
	Where <Xx> defines the language installed.

Solaris HSM client environment requirements

Before you install the HSM client, your workstation must meet minimum requirements.

Hardware requirements

- Machine
 - SPARCstation based on sun4u architecture or compatible workstation
- 100 MB of available disk space
- 200 MB of temporary disk space in the /var file system.
- 512 MB memory

Software requirements

- Sun Solaris 10 in 32-bit or 64-bit kernel mode (SPARC only)
- Veritas file system (VxFS) 4.1 or 5.0
- Tivoli Storage Manager Backup-Archive Client installed

Installing the Solaris HSM client (SPARC only)

Before installing the HSM client in a Solaris SPARC environment, you need to review the general requirements, system specific requirements, and review all steps. For your convenience the steps for installation the backup-archive client and API are included.

Follow these steps for a first installation versus a replacement installation:

1. Login as the root user and mount the 'Tivoli Storage Manager UNIX Client CD-ROM' to /cdrom.
2. Change to the directory where the packages are stored, for example: `cd /cdrom/tsmcli/solaris` (see Table 10 on page 18).

Note: If the 'Tivoli Storage Manager UNIX Client CD-ROM' is not mounted to /cdrom or if the packages are stored in a different directory (for example, if they were downloaded by ftp), change to the correct directory.

3. Issue the following commands to install the required API and the Tivoli Storage Manager Backup-Archive Client packages:

```
pkgadd -a ./tsmadmin -d ./TIVsmCapi.pkg TIVsmCapi
pkgadd -a ./tsmadmin -d ./TIVsmCba.pkg TIVsmCba
```

Note: These commands install the API shared libraries and samples; the backup-archive command line, GUI, Web client, and the administrative client.

4. Issue the following command to install the HSM package which includes the client commands, the Space Management Agent (hsmagent), and the HSM GUI:

```
pkgadd -a ./tsmadmin -d ./TIVsmChsm.pkg TIVsmChsm
```

Note: This command installs the HSM client package in the Global Zone only. The -G option of the **pkgadd** command is not required and is ignored if used. See "Solaris local zones support" on page 20 for more information on using the HSM client on Local Zones.

Solaris local zones support

On Solaris SPARC platforms you can add Space Management to the Veritas File Systems assigned to the Local Zones. Solaris x86 platforms are not supported.

You can only install the HSM client in the Global Zone. You cannot install the HSM client on a Local Zone. Only the administrator of the Global Zone has the privilege to add Space Management to the Veritas File Systems assigned to the Local Zones and this cannot be performed from the Local Zone itself.

No HSM client commands or HSM daemons is present in the Local Zones. The HSM client daemons that run in the Global Zone are responsible for the Local Zones. For example, if a migrated file belongs to a managed file system of the Local Zone it is automatically recalled (if required) by the recall daemon running in the Global Zone. A manual recall of migrated files is possible, but only by the administrator of the Global Zone. No special options or documentation is needed to use the HSM client on Local Zones (use Sun's documentation if you need information on administering zones).

See the "Installing the Solaris HSM client (SPARC only)" on page 19 for information on installing the HSM client in the Global Zone. See Chapter 4, "Adding and configuring space management," on page 47 for information on the process of adding space management to a file system.

Upgrading the Solaris HSM client

To upgrade the HSM client on a Solaris system, you need to uninstall the packages related to the Tivoli Storage Manager before installing the new version of the HSM client. Mixed installations are not supported. Make sure that you uninstall the packages in the given order.

Note:

- The Tivoli Storage Manager installation administration file, `tsmadmin`, is used in place of the default administration file, `/var/sadm/install/admin`, so that you will not be prompted for `setuid`, `setgid`, or `superuser` permission during this process. If you want to use the default administration file, remove the `-a ./tsmadmin` option from the following commands and make sure that you answer *yes* to the questions about `setuid`, `setgid`, and `superuser` permission during installation.
- You cannot uninstall a single component of this package (for example, the Web GUI) without first uninstalling the complete package.
- All activity on the HSM-managed file systems should be avoided during software installation. Access to migrated files during the upgrade will be not possible (the process trying to read a stub would hang).
- To completely remove the HSM client, without reinstalling a new version, you *must* first remove HSM support from all managed file systems. See "Uninstalling the Solaris HSM client" on page 21.

Follow these steps to uninstall all of the packages related to Tivoli Storage Manager (HSM, including command line, GUI, Web GUI, and administrative client components):

1. If you have installed the HSM client and you want to uninstall it in order to install a new version, issue the following command as root user:

```
pkgrm -n -a ./tsmadmin TIVsmChsm
```

2. Issue the following command to uninstall the Tivoli Storage Manager Backup-Archive Client. This will uninstall all of the components of the backup-archive client (command line, GUI, Web GUI, and the administrative client).

```
pkgrm -n -a ./tsmadmin TIVsmCba
```

3. All additional language messages packages must be removed before removing the API package. Issue the following command as root user to remove them:

```
pkgrm -n -a ./tsmadmin TIVsmClFr TIVsmClDe TIVsmClIt TIVsmClPt \  
TIVsmClEs TIVsmClCs TIVsmClHu TIVsmClPl TIVsmClRu \  
TIVsmClJa TIVsmClKo TIVsmClSc TIVsmClTc TIVsmClBc
```

4. Issue the following command to uninstall the Tivoli Storage Manager API. You must remove the Tivoli Storage Manager Backup-Archive Client before you remove the API.

```
pkgrm -n -a ./tsmadmin TIVsmCapi
```

5. To install a new version of the HSM client, see “Installing the Solaris HSM client (SPARC only)” on page 19.

Uninstalling the Solaris HSM client

To completely remove the HSM client, without reinstalling a new version, you first remove HSM support from all managed file systems and then remove the HSM client package.

If you plan to upgrade the HSM client, rather than complete remove it, see “Upgrading the Solaris HSM client” on page 20.

To completely remove the HSM client follow these steps:

1. To remove the HSM support, issue the following command:

```
dsmmigfs remove
```

2. To uninstall the HSM client, issue the following command as root user:

```
pkgrm -n -a ./tsmadmin TIVsmChsm
```

Registering your workstation with a Tivoli Storage Manager server

After installing the HSM client, your system must be registered as a client node with a Tivoli Storage Manager server before you can request services from that server. Your Tivoli Storage Manager server administrator has set the registration either to closed or to open.

Open registration

Open registration permits the root user to register your workstation as a client node with the server. Your workstation must be registered before anyone can use Tivoli Storage Manager on that node.

When you use open registration:

- 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 Tivoli Storage Manager storage. Users can delete archived files that they own.

Note: Your Tivoli Storage Manager administrator can change these defaults at any time.

Follow these steps to register your workstation with the Tivoli Storage Manager server:

1. Start a session with the HSM GUI by issuing the **dsmsmj** command or 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:
 - The initial password that you want to use, if a password is required
 - Contact information, such as your name, user ID, and telephone number
3. To register your workstation with additional servers, issue the **dsmsmj** command or the **dsmc** command with the *servername* option for each server. For example, where *dsmserv* is the name of a server that you identified in your `dsm.sys` file, issue the command:

```
dsmsmj -servername=dsmserv
```
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 Tivoli Storage Manager administrator with information to register your workstation as a client node with the server.

If your enterprise uses closed registration, provide the following information to your administrator:

1. Your node name: the value that the **hostname** command returns or the node name that you specified 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 policy domain to which your client node belongs

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

- The permission for you to compress files before you send them to the server
- The permission for you to delete backup and archive data from Tivoli Storage Manager storage

Managing your password

If you installed the HSM client on your client node, and you require a Tivoli Storage Manager password, set the *passwordaccess* option to *generate* in your `dsm.sys` file. Set the *mailprog* option in your `dsm.sys` file to send you the password each time it generates a new one.

Tivoli Storage Manager 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 a Tivoli Storage Manager password.

The backup-archive client and the HSM client use the same password when both clients contact the same Tivoli Storage Manager server. One password is required for each Tivoli Storage Manager server that your client node contacts for services.

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

If you specify a default server on the *defaultserver* option in your *dsm.sys* file and you do not specify a migration server on the *migrateserver* option, 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 Tivoli Storage Manager password can only be changed from the command line. To change that password from the command line, issue the command (where *oldpw* is the old password and *newpw* is your new password):

```
dsmsetpw oldpw newpw
```

The HSM GUI

The HSM GUI (the IBM Tivoli Enterprise Space Management Console) is a Java GUI that allows you to manage multiple HSM client systems and monitor HSM client activities. The HSM client systems can be local or remote and are managed by connecting to the HSM agent (also known as the space management agent).

The HSM GUI works with all HSM UNIX and Linux clients and can be used on any HSM system where the HSM agent is running.

The HSM GUI can also be used on Windows systems (Windows XP, and Windows server) so you can remotely administer multiple HSM clients from a single point. Download the Windows package from <ftp://ftp.software.ibm.com/storage/tivoli-storage-management/>.

Before using HSM GUI, check that:

1. A Java Runtime Environment (JRE) 1.5.x or JRE 1.6.x (where x is 1 or later) is installed.
2. HSM client is installed and the **hsmagent** is started on any HSM client machine that will be managed through the HSM GUI. See “Verifying the HSM agent is running” on page 24 to verify the hsmagent is running.
3. The Tivoli Storage Manager client node used to manage HSM is already registered and authenticated with the Tivoli Storage Manager Server. The node name and password required to sign on to an HSM node are identical to the node name used to manage HSM on the node and the password used to authenticate with the Tivoli Storage Manager server.

For more information on using the HSM GUI, access its online help from the **Help** menu.

Verifying the HSM agent is running

In order to use the HSM GUI to manage an HSM system, the HSM agent must be running.

Follow these steps to verify the HSM agent (space management agent) is running:

1. Log in or telnet to the HSM system that you want to manage with the HSM GUI by issuing the command:

```
telnet <HSM_machine_name>
```

2. Verify that the HSM client version installed matches the version of the HSM GUI. For example, from Solaris issue the command:

```
pkginfo -1 TIVsmChsm
```

3. Issue the **cat** command to verify that the *passwordaccess generate* option is used in the *dsm.sys* configuration file.

Note: If the *DSM_DIR* environment variable is not set, the *dsm.sys* configuration file in the installation directory is used.

```
cat $DSM_DIR/dsm.sys
```

4. Check that the authentication of the Tivoli Storage Manager server is enabled by issuing the command.

```
dsmadm q status
```

5. Ensure that the IBM Tivoli Storage Manager password is stored locally by issuing the following command. If the password is not found locally, you will be prompted to set it.

```
dsmc q sess
```

6. Ensure that the HSM agent (*hsmagent*) is running by issuing the following command.

```
ps -ef | grep hsmagent | grep -v grep
```

Note:

- If you have set the *DSM_DIR* environment variable, create a link or copy the *hsmagent.opt* XML configuration file located in the *hsm* installation directory to the *\$DSM_DIR* directory. If you do not perform this step, the HSM agent will not start correctly. For example issue the command:

```
ln -s /opt/tivoli/tsm/client/hsm/bin/hsmagent.opt $DSM_DIR/hsmagent.opt
```

- You can change HSM agent settings, such as the default port number *1555*, by editing the *hsmagent.opt* XML configuration file.

7. If the HSM agent is not running (no *hsmagent* process is found), start the HSM agent by issuing one of the following command sets:

```
cd /opt/tivoli/tsm/client/hsm/bin  
hsmagent
```

For AIX:

```
cd /usr/tivoli/tsm/client/hsm/bin  
hsmagent
```

Using the HSM GUI

When you use the HSM GUI to connect to a Tivoli Storage Manager client node, you can view and managed displayed file systems.

Follow these steps to start and use the HSM GUI:

1. On Windows, click the **Start** button and select **Programs** → **Tivoli Storage Manager** → **Space Management Console**. On UNIX or Linux run the **dsmsmj** command from a UNIX or Linux system where the HSM client is supported and installed. The Welcome Page opens.
2. On the left-hand side, select **Manage Resources** to connect to your preferred Tivoli Storage Manager client node or to create a list of your preferred Tivoli Storage Manager client nodes. When you log in to a client node, you must use the node name and password that were used to authenticate with the server. Check the **Save password locally** check box to save the password on the local machine. You must authenticate with at least one client node in order to retrieve the saved password.
3. On the right-hand side, select the root of the tree to display the list of preferred client nodes with their properties. A new client node can easily be added, edited, or removed by clicking on the related button.
4. Select a client node in the tree to display a table of file systems on the selected HSM node, including their different attributes.
5. Select a file system in the table to perform the following:
 - Add Space Management
 - Remove Space Management
 - Activate Space Management
 - Deactivate Space Management
 - Modify specific file system properties
6. Click on the related button to display a graphical overview of the file systems on the selected client node.

Chapter 3. Selecting HSM client options

You need to review your environment and each space management option so you can then set options in your `dsm.sys` and `dsm.opt` files. As optional tasks, you can assign management classes to your files and modify your include-exclude file.

Note: For more information about these options, see Chapter 11, “HSM client `dsm.sys` and `dsm.opt` option reference,” on page 89

These options determine such things as:

- The Tivoli Storage Manager server to which your files migrate or premigrate
- 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 must elapse after a migrated or premigrated file is deleted or modified on your local file system before the copy in Tivoli Storage Manager storage expires and is removed

During installation of the IBM Tivoli Storage Manager for UNIX 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/bin`
- On Solaris, HP-UX, and Linux x86/x86_64: `/opt/tivoli/tsm/client/ba/bin`

If you are installing both the backup-archive client and the HSM client at the same time, copy and rename the sample options files and 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 HSM client, you need to restart all space management daemons to activate the changes. All HSM-managed file systems must be mounted either automatically or manually each time you restart your system to enable space management.

For information about setting other Tivoli Storage Manager options, see IBM Tivoli Storage Manager for UNIX and Linux Backup-Archive Clients Installation and User’s Guide.

The HSM client shares the following common files and code with the IBM Tivoli Storage Manager for UNIX and Linux Backup-Archive Clients:

- Communication protocols
- `dsm.opt` file
- `dsm.sys` file
- Include-exclude file
- Node registration
- Server file spaces

Displaying HSM client options

You can display information about HSM client options either from the HSM GUI or the command line.

To display information about options from the command line, use the **dsmmigquery** command:

```
dsmmigquery -options
```

To display options from the HSM GUI, follow these steps:

Note:

You can use the HSM GUI's online help to get more detailed information by selecting **Help Topics** from the **Help** menu, by pressing the **F1** key, or by clicking the **?** icon in the upper right of the HSM GUI window.

1. From the **Manage Resource** task, select the client node for which you want to display client and server preferences.
2. Sign on to the selected client node if you are not already connected.
3. Click the **Client Node Properties** button or choose **View** → **Client Node Properties** from the menu.

Editing the options files using the backup-archive client

The HSM client shares the option files, `dsm.opt` and `dsm.sys`, with the IBM Tivoli Storage Manager for UNIX and Linux Backup-Archive Clients. To edit the options files you can use the Preferences Editor window of the backup-archive GUI. You can also edit the files in a text editor.

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.
- For Solaris operating systems, `dsm.sys` and `dsm.opt` are symbolic links to the actual files that you store in `/usr/bin`. This protects your files from deletion if you uninstall Tivoli Storage Manager.
- For more information on the options in the `dsm.opt` file, see “Editing `dsm.opt` options” on page 32.
- For more information on the options in the `dsm.sys` file, see “Editing `dsm.sys` options” on page 29.

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

1. Start the backup-archive GUI by issuing the command **dsmj**.
2. Open the Preferences Editor 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 dsm.sys options

The options you define in the dsm.sys file for the HSM client affect automatic migration, reconciliation, and recall.

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

In your 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 using the options parameters
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

Table 11 on page 30 provides a brief description of each space management option that you can set in your dsm.sys file.

Note:

- For information about setting up space management for AIX in an HACMP environment, see “Setting up space management for HACMP on AIX JFS2” on page 39.
- **For AIX GPFS and Linux x86/x86_64 GPFS file systems only:** To access your data on the server in the event of failover, ensure that all HSM nodes participating in the failover environment share the same node name or *asnodename* in your dsm.sys file.
- With client node proxy support it is possible to share the filespaces of one Tivoli Storage Manager node on the Tivoli Storage Manager server with other Tivoli Storage Manager nodes.
- With the exception of *errorprog*, place all space management options *at the beginning* of your dsm.sys file *before* any server stanzas that you define.
- You can specify both a default server and a migration server in your dsm.sys file. If you do not specify a server name with either the *defaultserver* option or the *migrateserver* option, then the server that you do specify in the first stanza of your dsm.sys file becomes your default server.
- If you specify a migration server with the *migrateserver* option, it overrides the server that you specify with the *defaultserver* option.

The following example displays an example of a dsm.sys server stanza containing options for a server that you want to contact.

```
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
KERNeImessages	no
OVERLAPRECALL	no
Servername	server1
COMMmethod	TCPip
TCPPort	1500
TCPSeveraddress	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
PARTnerluname	raptor
TPname	appcdel
CPICMOdename	appc
Passwordaccess	generate
Mailprog	/usr/bin/xsend root
Groups	system tsm
Users	sullivan tang stewart
Inclxcl	/adm/tsm/migrate.excl
ERRORProg	/bin/cat

Table 11. Space Management Options in Your Options Files

Option	Default	Description
candidatesinterval	1	Specifies how often the dsmscoutd daemon scans file systems for file information in those file systems which were added for space management. The range of values is 0 through 9999. See "Candidatesinterval" on page 89.
checkfororphans	no	Specifies whether or not the dsmreconcile command checks for orphans. Specify <i>yes</i> or <i>no</i> . See "Checkfororphans" on page 90.
checkthresholds	5	Determines how frequently the space monitor daemon checks space usage on the file systems to which you add space management. The range of values is 1 through 9999. See "Checkthresholds" on page 90.
compression	no	Compresses files before you send them to the Tivoli Storage Manager server. Specify <i>yes</i> or <i>no</i> . See "Compression" on page 91.
defaultserver		Specifies the name of the default Tivoli Storage Manager server to contact for space management services if you do not specify a server name on the <i>migrateserver</i> option. See "Defaultserver" on page 91.
errorprog		Specifies a program to which you want to send a message if a severe error occurs during space management processing. Place this option <i>after</i> all space management options for each server that you define in your dsm.sys file. See "Errorprog" on page 92.
filelist		Use with Tivoli Storage Manager for Space Management shell commands to process a list of files. See "Filelist" on page 92.
inclxcl		Specifies the file name and path of your include-exclude file. See "Modifying the include-exclude options file" on page 35. See "Inclxcl" on page 93.
maxcandprocs	5	Specifies the number of parallel processes in the scout daemons which can scan file systems. The range of values is 2 through 20. See "Maxcandprocs" on page 94.

Table 11. Space Management Options in Your Options Files (continued)

Option	Default	Description
maxmigrators	5	Specifies the number of parallel migration sessions that you can perform. The range of values is 1 through 20. If this option is changed from the default, a corresponding increase should be made in the Tivoli Storage Manager server configuration to update the hsm node <i>MAXNUMMP</i> value. See “Maxmigrators” on page 94.
maxrecalldaemons	20	Specifies the maximum number of recall daemons that you can run at one time to perform recalls for your client node. The range of values is 2 through 99. If <i>maxrecalldaemons</i> value is less than <i>minrecalldaemons</i> value, the default is set for both options. See “Maxrecalldaemons” on page 94.
maxreconcileproc	3	Specifies the maximum number of automatic reconciliation processes that you can start at one time. The range of values is 1 through 99. See “Maxreconcileproc” on page 95.
maxthresholdproc	3	Specifies the maximum number of automatic threshold migration processes that you can start at one time. The range of values is 1 through 99. See “Maxthresholdproc” on page 95.
migfileexpiration	7	Specifies the number of days that copies of migrated or premigrated files remain on the Tivoli Storage Manager server after they are modified on, or erased from, your local file system. The range of values is 0 through 9999. See “Migfileexpiration” on page 96.
migrateserver		Specifies the Tivoli Storage Manager server to which you want to migrate files from your client node. Specify only one migration server for each client node. See “Migrateserver” on page 96.
minmigfilesize	0	Specifies the minimum file size for a file to be eligible for migration. The range of values is 0 through 2147483647. See “Minmigfilesize” on page 97.
minrecalldaemons	3	Specifies the minimum number of recall daemons that you want to run at the same time to perform recalls for your client node. The range of values is 1 through 99. See “Minrecalldaemons” on page 97.
optionformat	<i>standard</i>	Specifies the format to use (either standard or short) when you specify HSM client commands. If you change the value, the new value is effective immediately. Select <i>standard</i> or <i>short</i> format. See “Optionformat” on page 98.
reconcileinterval	24	Specifies how often automatic reconciliation of file systems occurs to which you add space management. The range of values is 0 through 9999. See “Reconcileinterval” on page 98.
restoremigstate	<i>yes</i>	Restores a file if it is backed up after migration. You cannot restore a stub file if you back up a file before migration because a server stub file copy does not exist. Specify <i>yes</i> or <i>no</i> . See “Restoremigstate” on page 99.

1. Obtain the server information from your Tivoli Storage Manager 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 to which 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.

Editing dsm.opt options

Most of the options in the `dsm.opt` file affect the backup-archive client. However two options `optionformat` and `restoremigstate`, affect both the backup-archive client and the HSM client.

Table 12 provides a brief description of the space management options that you can set in your `dsm.opt` file.

Note: The Tivoli Storage Manager `nfstimeout` option is ignored for HSM clients the HSM client only operates on local file systems.

To edit the space management options in your `dsm.opt` file:

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 12. HSM client options in your `dsm.opt` file

Option	Default	Description
<code>defaultserver</code>	The server that you identify in the first stanza of your <code>dsm.sys</code> file.	Specifies the name of the default Tivoli Storage Manager server to contact for space management services if you do not specify a server name on the <code>migrateserver</code> option. The value of <code>defaultserver</code> in <code>dsm.sys</code> overrides <code>defaultserver</code> in <code>dsm.opt</code> . See “Defaultserver” on page 91.
<code>migrateserver</code>	The value of <code>defaultserver</code> option.	Specifies the Tivoli Storage Manager server to which you want to migrate files from your client node. Specify only one migration server for each client node. The value of <code>migrateserver</code> in <code>dsm.sys</code> overrides <code>migrateserver</code> in <code>dsm.opt</code> . See “Migrateserver” on page 96.
<code>optionformat</code>	Standard	Specifies the format to use (either standard or short) when you specify HSM client commands. If you change the value, the new value is effective immediately.
<code>restoremigstate</code>	Yes	Restores a file to stubbed (migrated) state if it is backed up after migration or premigration. Tivoli Storage Manager records the migration state of files during backup, so only those files that were migrated or premigrated at the time of backup can be restored to stubbed (migrated) state.

3. Set the `DSM_CONFIG` environment variable to point to your `dsm.opt` file. For instructions to set this variable, see “Setting environment variables” on page 37.

Optional setup features

To facilitate using the HSM client, you can create management classes, create include and exclude options, and set environment variables.

Management classes that you assign to files determine whether or not a file is eligible for backup or migration. An include-exclude options file contains statements that identify any files you want to include or exclude from backup or migration. Environment variables point to files that the Tivoli Storage Manager uses.

Assigning management classes to files

Your Tivoli Storage Manager 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 using 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. For information about backup and archive copy groups, see the Tivoli Storage Manager for UNIX Backup-Archive Clients Installation and User's Guide. Table 13 lists the space management attributes and their defaults that might be included in a management class.

Table 13. HSM client attributes in a management class

Attribute	Default	Description
spacemgtechnique	<i>None</i>	Specifies that a file is eligible for automatic and selective migration, selective migration only, or neither. The values for this attribute are: <i>Auto</i> The file is eligible for both automatic and selective migration. <i>Selective</i> The file is eligible for selective migration only. <i>None</i> The file is not eligible for migration. Note: If you use the default management class named Standard that is shipped with the Tivoli Storage Manager, and your administrator has not changed the default setting for the <i>spacemgtechnique</i> attribute, files are not migrated from your workstation.
automignonuse	<i>0</i>	Specifies the number of days (<i>0</i> through <i>9999</i>) that must elapse since you last accessed the file before it is eligible for automatic migration.
migrequiresbkup	<i>yes</i>	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 values for this attribute are: <i>yes</i> , a current backup version must exist. <i>no</i> , a current backup version is optional. Note: If you set this attribute to Yes in the management class that you assigned to a file, Tivoli Storage Manager checks for a current backup version of the file on your migration server only. If a current backup version <i>does not exist</i> on your migration server, the file is not migrated, even if a current backup version exists on another server.
migdestination	spacemgpool	Specifies the name of the storage pool in which Tivoli Storage Manager stores the file when it migrates.

Table 14 defines some tasks you might perform and the type of management class to assign to your files.

Table 14. 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.

Table 14. Assigning management classes to your files (continued)

Task	Assignment
Your client node communicates with the same server for both space management and backup-archive services.	Assign a management class to a file containing space management attributes <i>and</i> backup-archive copy groups that you assigned to Tivoli Storage Manager 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 <code>/home/holland</code> to Server1, the include-exclude options file that you use for Server1 might specify a management class named <code>mgmt1a</code> for a file named <code>/home/holland/testfile</code>. This management class must contain an appropriate backup copy group for the file.</p> <p>If you migrate files in the <code>/home</code> file system to Server2, the include-exclude options file that you use for that server might specify a management class named <code>mgmt2b</code> for the same file. That management class must contain appropriate space management attributes for the file.</p>

Displaying management class information

You can display management class information from the HSM GUI for the default migration server only. If your client node contacts one or more additional servers for backup and archive services, use the backup-archive client GUI or the **dsmmigquery** command to display information about available management classes on those servers.

For more information about management classes or include-exclude options, see IBM Tivoli Storage Manager for UNIX and Linux Backup-Archive Clients Installation and User's Guide.

To display information from the command line about management classes that you can assign to your files, issue the command:

```
dsmmigquery -mgmtclass -detail
```

To use the HSM GUI to display policy and management classes that you can assign to your files, follow these steps:

Note: To access the HSM GUI online help select the **Help Topics** from the **Help** menu, press **F1**, or click the **?** icon in the upper right of the HSM GUI window.

1. From the **Manage Resource** task, select the client node for which you want to display policy information.
2. Sign on to the selected client node, if not already connected.
3. Click the **Client Node Properties** button or choose **View** → **Client Node Properties** from the menu.
4. Click on the **Policy** tab at the left side of the Client Node Properties window.
5. Select a management class and click **View Details** to display the attributes of the management class.

Modifying the include-exclude options file

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, such as system files or files that the HSM client creates and uses, on your local file system at all times, 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 AIX GPFS and Linux x86/x86_64 GPFS file systems do not use the *inlexcl* 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 participating in the local failover group or node set.

For all other file systems, use the *inlexcl* 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 Tivoli Storage Manager server that your client node contacts for services. For example, if your *dsm.sys* file contains two stanzas with options for two servers, you can include an *inlexcl* option in each stanza. Each *inlexcl* 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 use the options in to include or exclude files from either backup or migration:

- Use include statements include files for backup or migration.
- Use exclude statements 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.

Table 15. Include and Exclude Statements

Option	Description
<i>exclude</i>	Excludes a file or a group of files from space management.
<i>exclude.backup</i>	Excludes a file from backup.
<i>exclude.file</i>	Excludes a file or a group of files from space management.
<i>exclude.file.spacemgmt</i>	Excludes a file from HSM services only. Use this option when you have both the backup-archive client <i>and</i> the HSM client installed.
<i>exclude.spacemgmt</i>	Excludes files and directories from HSM services only. Use this option when you have both the backup-archive client <i>and</i> the HSM client installed.
<i>include</i>	Includes files for backup and HSM services or LAN-free data transfer.
<i>include.file</i>	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, and the files that you included in that list migrate, each file is tested against the options in the include-exclude options file *beginning with the last option that you specified and working up*. 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 migrating to storage from a local file system, assign the file to a management class without automatic or selective migration.

Creating an include-exclude list

The HSM client shares the include-exclude list with the IBM Tivoli Storage Manager for UNIX and Linux Backup-Archive clients. You can use the backup-archive GUI or the command line to create and edit your include-exclude list.

Note:

- Before using 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** button on the Preferences Editor window.

To create or edit the include-exclude list using the client backup-archive GUI, follow these steps:

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 the **Add** button 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.

To create an include-exclude list from the command line:

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 your *dsm.sys* file, use the *inclxcl* option to define the name of the include-exclude options file you created in 1.

Setting environment variables

Set your environment variables to point to files that Tivoli Storage Manager uses.

Note:

- You cannot specify the root directory for DSM_DIR, DSM_CONFIG, or DSM_LOG.
- Use the LANG environment variable to specify the language that you want to use.

Table 16. Environment Variables

Variable	Description
DSM_DIR	Points to the resource files, the dsm.sys file, and the executable file, dsmtca. If you do not set DSM_DIR, the HSM client searches for the executable files in the installation directory.
DSM_CONFIG	Points to your dsm.opt file. <ul style="list-style-type: none">• If you do not set DSM_CONFIG, the HSM client searches for the options file in the directory to which DSM_DIR points.• If you do not set DSM_DIR, the HSM client searches for the options file in the installation directory.• For Solaris operating systems, dsm.sys and dsm.opt are symbolic links to the actual files that you store in /usr/bin. This protects your files from deletion if you uninstall Tivoli Storage Manager.
DSM_LOG	Points to the directory where you want the dsmerror.log file to reside. The error log file contains information about any errors that occur during processing. This log file helps Tivoli Customer Service diagnose severe errors. <ul style="list-style-type: none">• If you define DSM_DIR but you do not define DSM_LOG, messages are written to dsmerror.log in the directory that you specified in DSM_DIR.• If you do not define DSM_LOG or DSM_DIR, error messages are written to dsmerror.log in the current directory. You receive a warning message if Tivoli Storage Manager cannot write messages to the log file. Processing continues.

Setting Bourne and Korn shell variables

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

For example:

```
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 user's directory and the path for the executable files, the resource files, and the dsm.sys file.
- The *DSM_CONFIG=/home/hsmuser/dsm.opt* entry identifies the path and file name for the dsm.opt file.
- The *DSM_LOG=/home/hsmuser* entry identifies the directory where you want the dsmerror.log file to reside.

Setting C shell variables

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

For example:

```
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 dsm.opt file. The /home/hsmuser path identifies the directory where you want to store the dsmerror.log file.

Setting up LAN-free data transfer for HSM

The HSM client on AIX GPFS, AIX JFS2, and Linux GPFS clients 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 Tivoli Storage Manager 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 Tivoli Storage Manager Storage Agent. The Tivoli Storage Manager Storage Agent must be installed on the same system as the client.

LAN-free prerequisites

Before establishing LAN-free support in your HSM environment, you need to check the prerequisites.

To enable LAN-free support:

- A Tivoli Storage Manager Version 5.5 or later client and server is required.
- You must install and configure the storage agent on the client workstation. For more information, refer to the following publications:
 - IBM Tivoli Storage Manager for AIX Storage Agent User's Guide, SC32-0129
 - IBM Tivoli Storage Manager for Linux Storage Agent User's Guide, SC32-0131

LAN-free options

After installing and configuring the IBM Tivoli Storage Manager Managed System for SAN feature on the client workstation, use options to enable LAN-Free data transfer.

Table 17 lists options you can use to setup LAN-free data transfer. See the IBM Tivoli Storage Manager for UNIX and Linux Backup-Archive Clients Installation and User's Guide for more information about these options.

Table 17. LAN-free data transfer options

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

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

To include files: Assuming that /hsm1 is an HSM 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 command:

```
include /hsm1/clientdata/lanfree lanfreemgmtclass
```

To exclude files: Assuming that /hsm1 is an HSM 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 command:

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

Setting up space management for HACMP on AIX JFS2

The HSM client permits your HSM-managed file system to become a part of a resource group in an HACMP cluster so that if there is a system failure you can obtain access to your data from another system. The HACMP environment is valid on AIX JFS2 file systems only and supports cascading and rotating takeover relationships in the same manner as the backup-archive client.

The password handling is the same as that of the HACMP backup-archive client. The dsm.opt file and the dsm.sys file can reside in the standard /usr/tivoli directory, or you can use the DSM_CONFIG environment variable that you set in the start script. However, the behavior must be identical on all nodes.

To set up space management in an HACMP environment you must configure HSM on every cluster node that manages or will take over HSM file systems from other nodes.

The HACMP setup in conjunction with HSM is supported at the following levels:

- HACMP 5.4 on AIX 5.3 with Tivoli Storage Manager HSM JFS2 6.1 client
- HACMP 5.4 on AIX 6.1 with Tivoli Storage Manager HSM JFS2 6.1 client
- HACMP 5.5 on AIX 6.1 with Tivoli Storage Manager HSM JFS2 6.1 client

For more information about the HACMP environment for AIX, see IBM Tivoli Storage Manager for UNIX and Linux Backup-Archive Clients Installation and User's Guide.

Across different nodes, the server stanza must have the following characteristics:

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


```

`kill -15 $pid 2>&1 1>> $LOG`;
`proccwait -v $pid 2>&1 1>> $LOG`;

$pid = `ps -aef|awk '/dsmrootd/{if ( \ $3 == "1" ) print \ $2}'`;
`echo "Killing dsmrootd ... $pid" >> $LOG`;
`kill -15 $pid 2>&1 1>> $LOG`;
`proccwait -v $pid 2>&1 1>> $LOG`;

# killing hsmagent
$pid = `ps -aef| awk '/hsmagent/{if ( \ $3 == "1" ) print \ $2}'`;
`echo "Killing hsmagent ... $pid" >> $LOG`;
`kill -15 $pid 2>&1 1>> $LOG`;
`proccwait -v $pid 2>&1 1>> $LOG`;

# starting the daemons with the right dsm.sys and dsm.opt
`echo "Starting dsmonitor d ..." >> $LOG`;
`echo \ $DSM_DIR \ $DSM_CONFIG >> $LOG;
dsmonitor d 2>&1 1>> $LOG`;
`echo "Starting dsmreca l d ..." >> $LOG`;
`echo \ $DSM_DIR \ $DSM_CONFIG >> $LOG;
dsmreca l d 2>&1 1>> $LOG`;
`echo "Starting dsmscou t d ..." >> $LOG`;
`echo \ $DSM_DIR \ $DSM_CONFIG >> $LOG;
dsmscou t d 2>&1 1>> $LOG`;
`echo "Starting dsmroot d ..." >> $LOG`;
`echo \ $DSM_DIR \ $DSM_CONFIG >> $LOG;
dsmroot d 2>&1 1>> $LOG`;

# starting hsmagent
`echo "Starting hsmagent ..." >> $LOG`;
`echo \ $DSM_DIR \ $DSM_CONFIG >> $LOG;
hsmagent 2>&1 1>> $LOG`;

# transfer the command line into an array
while(my $temp=shift @ARGV){
    $FS[@FS]=$temp;
}
# Reactivate all filesystems
`echo "Starting to reactivate FS.." >> $LOG`;
for(my $i=0; $i < @FS; $i++){
    `dsmmigfs reactivate $FS[$i] 2>&1 1>> $LOG`;
}

`echo "HSM is Started" >> $LOG`;

```

Figure 2. stop_HSM : Stops HSM

```

#! /usr/bin/perl
#-----
# call with
# stop_HSM [filesystems]
#-----

my $LOG="/tmp/HSM.log"; # place of your logfile

$ENV{"DSM_DIR"}="[your location of dsm.sys]";
$ENV{"DSM_CONFIG"}="[your location of option file]";
`echo "*****" >> $LOG`;
`echo "* Stopping HSM *" >> $LOG`;
`echo "*****" >> $LOG`;

`echo STOP\\t\\t\\t \ `date` >> $LOG`;
`echo \ $DSM_DIR \ $DSM_CONFIG >> $LOG`;

# Stopping all daemons

```

```

my $pid = `ps -aef|awk '/dsmrecalld/{if ( \ $3 == "1" ) print \ $2}'`;
`echo "Killing dsmrecalld ... $pid" >> $LOG`;
`kill -15 $pid 2>&1 1>> $LOG`;
$pid = `ps -aef|awk '/dsmmonitord/{if ( \ $3 == "1" ) print \ $2}'`;
`echo "Killing dsmmonitord ... $pid" >> $LOG`;
`kill -15 $pid 2>&1 1>> $LOG`;
$pid = `ps -aef|awk '/dmscoutd/{if ( \ $3 == "1" ) print \ $2}'`;
`echo "Killing dmscoutd ... $pid" >> $LOG`;
`kill -15 $pid 2>&1 1>> $LOG`;
$pid = `ps -aef|awk '/dsmrootd/{if ( \ $3 == "1" ) print \ $2}'`;
`echo "Killing dsmrootd ... $pid" >> $LOG`;
`kill -15 $pid 2>&1 1>> $LOG`;

# killing hsmagent
$pid = `ps -aef| awk '/hsmagent/{if ( \ $3 == "1" ) print \ $2}'`;
`echo "Killing hsmagent ... $pid" >> $LOG`;
`kill -15 $pid 2>&1 1>> $LOG`;

# transfer the command line into an array
while(my $temp=shift @ARGV){
    $FS[@FS]=$temp;
}

# Stop all processes using this filesystem
`echo "Killing the rest ..." >> $LOG`;
for(my $i=0; $i < @FS; $i++){
    @PROCS=split (" ",`fuser -c $FS[$i] 2>/dev/null`);
    while(my $PROC=shift @PROCS){ $PROC=~s/c//g;
        `kill -15 $PROC 2>&1 1>> $LOG`;
    }
}

# now you can deactivate the filesystems
`echo "Starting to deactivate FS.." >> $LOG`;
for(my $i=0; $i < @FS; $i++){
    `dsmmigfs deactivate $FS[$i] 2>&1 1>> $LOG`;
}

`echo "HSM is stopped" >> $LOG`;

```

Space management for AIX GPFS and Linux x86/x86_64 GPFS clusters

The HSM client permits your HSM-managed file system to become a part of a resource group in the cluster so that if there is a system failure, you can obtain access to your data from another system.

The HSM client for AIX GPFS and Linux x86/x86_64 GPFS 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 can reside in the standard `/opt/tivoli` directory for Linux x86/x86_64 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.

To set up space management in an HACMP environment you must configure HSM on every cluster node that manages or will take over HSM file systems from other nodes.

- Across different nodes, the server stanza must have the following characteristics:
 - The server stanza should contain a node name, and the node name should be the same for the server stanza on each system.
 - The node name can be a cluster name. It does not need 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 automigration and recall capabilities to other GPFS nodes within a cluster environment by enabling failover and using the *asnodename* option to share the GPFS filespace on the server. You must have distributed automigration and recall to enable failover and to share the GPFS filespace with the *asnodename* option. This is only supported on a Tivoli Storage Manager Version 5.3 or higher HSM client, Tivoli Storage Manager client, and Tivoli Storage Manager server.

To enable failover of automigration and recall capabilities follow these steps:

1. Install the HSM client on all participating GPFS nodes in the cluster environment (see “AIX GPFS HSM client installation overview” on page 7 or “Linux x86/x86_64 GPFS HSM client installation overview” on page 14).
2. Register each HSM client node with the Tivoli Storage Manager server. These are the nodes that are being used for managing HSM file systems and those that are being used for distributed HSM. Nodes that are only being used to access data of HSM-managed file systems do not need to have HSM installed, and they do not need to be registered on the Tivoli Storage Manager server.
3. Issue the **dsmmigfs enableFailover** command on the cluster node that is to be the owner node to enable failover of HSM management of the GPFS file systems.
4. The Tivoli Storage Manager server administrator must use the **grant proxynode** server command to grant proxy authority to the source nodes to access the node name where the GPFS file systems are stored.
5. Update the options file on each source node by specifying the *asnodename* option to access the common filespace for the HSM managed file systems on the Tivoli Storage Manager server.
6. Enable failover of HSM on all cluster nodes that will participate in distributed HSM (owner and source nodes) with the **dsmmigfs enableFailover** command.

Limitations of HSM support for AIX GPFS and Linux x86/x86_64 GPFS

Be aware of some limitations of HSM support for AIX GPFS and Linux x86/x86_64 GPFS systems.

The limitations are:

- 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. Those 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 AIX GPFS and Linux x86/x86_64 GPFS systems is not completely integrated with the backup-archive client support. For example, the HSM client refers to the `<fs>/SpaceMan/hsmfsconfig.xml` file to determine which server to use for a file system. The client might contact a different server for each file system. In contrast, the backup-archive client determines which server to use from the `dsm.opt` file, the `dsm.sys` file, or from a parameter that you specify on the command line when you start a backup-archive command line client. A backup-archive client process might back up, archive, restore, or retrieve from one

server. If you need backup-archive client services for different servers, start a new backup-archive client process.

Configuring the HSM space management agent

The HSM agent program, `hsmagent` (also called the space management agent), is the communication endpoint for the HSM GUI. After starting the HSM agent, the HSM GUI can connect to it on the port specified in the `hsmagent.opt` file.

The HSM agent can be started with the `hsmagent` command with no command-line parameters. On AIX GPFS and Linux x86/x86_64 GPFS the HSM agent is started from the watch daemon.

The password file must contain the password for the node that is running the HSM client to connect the HSM GUI with the HSM agent.

If you have set the `DSM_DIR` environment variable, the space management agent searches for the `hsmagent.opt` file in this location. On AIX GPFS the `hsmagent.opt` file is located in `/usr/tivoli/tsm/client/hsm/bin`. On Linux x86/x86_64 GPFS, HP-UX and Solaris the file is located in `/opt/tivoli/tsm/client/hsm/bin`.

The case-sensitive `hsmagent.opt` file is an XML configuration file. The default `hsmagent.opt` file, which is delivered with the product, looks similar to Figure 3 (Table 18 on page 45 describes the options):

Figure 3. XML configuration file

```
<?xml version='1.0' encoding='ISO-8859-1' ?>
<HSMAGENT>
  <Options>
    <!-- Portnumber -->
    <PortNumber type="int">1555</PortNumber>
    <!-- Timeout for waiting for connection. In msec
    A timeout of 0 disables the timeout check, default is 2000 -->
    <AgentTimeOut type="int">2000</AgentTimeOut>
    <!-- Timeout of a Session in sec -->
    <SessionTimeOut type="int">3600</SessionTimeOut>
    <!-- Location and file for the Gui Options -->
    <GuiOptionsFile type="string">hsmagent.guioptions</GuiOptionsFile>

    <Tracing>
    <!-- Location and name of tracefile. Is empty and
    traceflags are set, output is consol -->
    <!-- Use
    <TraceFile type="string">./MyTrace.txt</TraceFile>
    Tracefile in local directory -->
    <TraceFile type="string"></TraceFile>
    <!-- List of traceflags
    Currently ENTER, EXIT, ERROR, COMM, SM are supported -->
    <TraceFlags type="list">
    <!--Use
    <Flag type="string">ENTER</Flag>
    to activate a trace flag -->
    </TraceFlags>
  </Tracing>
</Options>
</HSMAGENT>
```

Table 18. HSM agent options

Parameter	Description
<i>PortNumber</i>	The <i>PortNumber</i> option specifies on which port the Space Management GUI can connect to the space management agent. Valid values are 1500 to 32152.
<i>AgentTimeOut</i>	The <i>AgentTimeOut</i> option should not be changed. You should only change the parameter if a service representative asks you to do so.
<i>SessionTimeOut</i>	The <i>SessionTimeOut</i> option should not be changed. You should only change the parameter if a service representative asks you to do so.
<i>GuiOptionsFile</i>	The <i>GuiOptionsFile</i> option is currently not used.
<i>Tracing</i>	The <i>Tracing</i> , <i>TraceFile</i> and <i>TraceFlags</i> option should only be used if a service representative asks you to do so.

Chapter 4. Adding and configuring space management

You need to explicitly add space management to a file system you want managed under HSM. Before you do, however, you need to understand concepts and steps involved in adding and configuring space management on your file systems. You can also deactivate, reactivate, and remove space management.

Note:

- To perform space management tasks, you must have root user authority.
- The HSM client can only work on mounted file systems. Before you can migrate files to Tivoli Storage Manager storage, first mount your file systems either automatically or manually and then add space management.
- During the mount process and while the HSM client is adding space management to your file systems, do not attempt to access any files in your file systems or perform any tasks against your file systems.

When you add space management to a file system, the HSM client:

- Creates a hidden directory for the file system named `.SpaceMan` that stores certain information objects required to perform space management tasks. See Appendix A, “The SpaceMan directory,” on page 135 for more information.
- 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 Tivoli Storage Manager backup-archive client. Copy the `hsmfsconfig.xml` to a directory that is included for automatic backups so that the Tivoli Storage Manager backup-archive client can process the file.

- Updates mount information for the native file system.

For Solaris VxFS file systems: If the `/etc/vfstab` file contains an entry for your file system, a new attribute is not added.

- Starts space management for your file system.

For AIX GPFS and Linux x86/x86_64 GPFS file systems: The `dsmwatchd` daemon starts at system startup (`inittab` entry). Ensure that DMAPI is enabled on all GPFS file systems that the HSM client manages. Issue the following command to query this information:

```
/usr/lpp/mmfs/bin/mmfsfs <DevicePath> -z
```

If DMAPI is disabled, enable it with following command:

```
/usr/lpp/mmfs/bin/mmchfs <DevicePath> -z yes
```

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

The AIX or Linux x86/x86_64 cluster node to which you add a GPFS file system becomes the preferred node for your file system. If several HSM clients on several AIX or Linux x86/x86_64 cluster nodes are candidates for managing a GPFS file system, the preferred node has precedence. In the event of failure and takeover by another node within the GPFS cluster, the HSM client that manages

the file system is not the HSM client on the preferred node. A GPFS cluster is a set of AIX or Linux x86/x86_64 cluster nodes that can mount a defined group of GPFS file systems.

HSM stub file considerations

HSM-created stub files on your space-managed file systems are bound to the space-managed file system itself. Because of this, you **cannot** perform the following:

- Move stub files as native stub files (without the migrated data) to other file systems, whether or not they are space-managed.
- Use the Tivoli Storage Manager raw device backup utility or any other image-based backup utility to restore a space-managed file system to another 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 additional considerations:

- You cannot add space management to file systems such as the root file system and the temp file system.
- To add space management to a nested file system, see “Adding space management to nested file systems” on page 49.
- To add space management to an exported file system, see “Adding space management to an exported file system” on page 49.

Recommendation: Do not add space management to the /usr and /var file systems. All of those file systems contain files that your operating system uses on a regular basis.

Adding space management to file systems

You can add space management by issuing a command or by using the HSM GUI.

To add space management to your local file systems from the command line, where, for example, the file system is /home the high threshold is 85 percent and the low threshold is 75 percent, issue the following command:

```
dsmmigfs add -ht=85 -lt=75 /home
```

For more information on this command see, “dsmmigfs add and update” on page 112.

To add space management to a file system using the HSM GUI, follow these steps:

1. From the HSM GUI Manage Resources window, click a client node in the **Client Nodes** tree. If the client node is not connected, then you need to sign on to the client node. The list of file systems on a client node appears in a table in the work area next to the **Client Nodes** tree.
2. Select a file system whose status is **Not Managed** from the work area, and click the **Manage** button. The status of the file system becomes **Active**.
3. If you want to change the default space management settings, click the **File System Properties** button to open the File System Properties window. Make and save any needed changes.

Adding space management to nested file systems

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.

To add space management to a parent file system and any nested file systems, perform all of the following steps. To add space management to a parent file system but not to its nested file system, perform steps one through three.

1. Unmount the nested file systems.
2. Add space management to the parent file system.
3. Remount the nested file systems.
4. Add space management to each nested file system. See “Adding space management to file systems” on page 48.

Mounting the parent file system before nested file system

For an AIX JFS nested file system that mounts automatically when you restart your system, you need to make sure you mount the parent system before mounting the nested file system.

Note: These steps are for AIX JFS2 file systems only.

Follow these steps to mount the parent file system before you mount the nested file system:

1. Issue the command `smit jfs`.
2. Select **Change/Show Characteristics of a Journaled File System**.
3. Select the nested file system.
4. Set **Mount AUTOMATICALLY at system restart?** to **no**.
5. Select **Do**.
6. Edit the `/etc/rc.adsmhsm` file that is shipped with the HSM client and add the following statement in the nested file system section at the end of the file (where `/test/migfs1` is the name of the nested file system):

```
mount /test/migfs1
```

Adding space management to an exported file system

HSM only supports NFS as a network file system interface to export a file system.

Note: CIFS implementations such as AIX FastConnect and SAMBA in their current forms do not enable sufficient concurrency for HSM environments and are not supported.

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

1. Instruct all NFS clients to unmount the exported file system.
2. To view which clients mounted the exported file system, issue:

```
/usr/sbin/showmount -a
```

3. Add space management to your file system using the **dsmmigfs add** command. See “Adding space management to file systems” on page 48 for more information.
4. To export the NFS file system again, issue:
`/usr/etc/exportfs -a`
5. Instruct all NFS clients to mount the exported NFS file system again.

Adding space management to WPARs on AIX 6.1

AIX 6.1 workload partitions (WPARs) act and look like a stand alone systems and provide an isolated environment for enterprise applications and data. An HSM client installed in the global partition has access to all file system data across WPARs,

Note: For more information on editing your `dsm.sys` file for WPARs, see IBM Tivoli Storage Manager for UNIX and Linux Backup-Archive Clients Installation and User’s Guide.

To add space management support for WPAR, 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 will be transparently recalled. They will also be transparently recalled if access from the global partition.

Updating settings from the command line

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

For example, to change the high threshold to `80` and the low threshold to `70` for the `/home` file system, issue the command:

```
dsmmigfs upd -ht=80 -l=70 /home
```

See “`dsmmigfs add and update`” on page 112.

Updating settings from the HSM GUI

You can update space management settings to change the way space usage is managed by HSM on an active or deactivated file system. If you update space management settings for a deactivated file system, the new settings do not take effect until you reactivate the file system.

Updating space management settings can affect the following:

- The order in which files are migrated to the Tivoli Storage Manager server during threshold migration and demand migration
- When threshold migration begins and ends
- The amount of space available on the server to store migrated and premigrated files

- The stub file size

To update space management settings for a file system, follow these steps:

1. From the Manage Resources window, click a client node in the **Client Nodes** tree. If the client node is not connected, then you need to sign on to the client node. The list of file systems on a client node displays in a table in the work area next to the **Client Nodes** tree.
2. Select a file system from the work area, and click the **File System Properties** button. The File System Properties window opens.
3. Update the space management settings in the Thresholds, Management, or Candidates tab in the File System Properties window.
4. Click **OK** to update the space management settings.

Space management settings

After you add space management to your file systems, you need to define space management settings to control different aspects.

Specifically, space management settings control:

- 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 Tivoli Storage Manager 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 HSM client maintains on your local file system
- The minimum size (in megabytes) that a file must have to qualify for partial file recall
- The minimum size (in megabytes) that a file must have to qualify for streaming recall mode

The information for each file system is stored in the `hsmfsconfig.xml` file, which is located in the appropriate `.SpaceMan` directory. You can modify the settings in the `hsmfsconfig.xml` file using the **dsmmigfs** command. You can also use the **dsmmigfs** command to display the settings in this file (see “`dsmmigfs query`” on page 117).

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 automigration and recall capabilities to source nodes within a cluster environment. This is only supported on a Tivoli Storage Manager Version 5.3 or higher HSM client, Tivoli Storage Manager client, and Tivoli Storage Manager server. See Chapter 3, “Selecting HSM client options,” on page 27 for more information.

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 migrate after you make the changes.
- If you reduce the quota, and the data that you currently migrate and premigrate 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.

Minimum migration file size

The HSM client does not migrate a file unless doing so saves space on your local file system. The exact minimum file size is dependant upon your file system; however, in general, the migrated file must be larger than the replacement stub file.

For AIX JFS2 file systems, before a file is considered for migration, the size must be greater than both the stub file size plus one byte and the file system fragment size. For example, if the stub file size is 2047 bytes, and the defined fragment size for a file system is 4096 bytes, the file size must be greater than 4096 bytes before it is eligible for migration.

For AIX GPFS and Linux x86/x86_64 GPFS file systems, the file size must be greater than the stub file size before it is eligible for migration.

For Solaris VxFS and HP-UX file systems, the file size must be greater than both the stub file size and the file system block size before it is eligible for migration.

You can use the *minmigfilesize* option to set the minimum file size considered for migration. If the *minmigfilesize* option is set, the size of a file must be greater than this option value before the file is eligible for migration. This option can be set either globally in the *dsm.sys* file or for a specific file system (using the **dsmmigfs** add and update commands). A valid file-system specific value will take precedence over the global *minmigfilesize* option setting specified in the *dsm.sys* file. See “*dsmmigfs* add and update” on page 112.

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 HSM client checks space usage on your file system, files automatically begin migrating to Tivoli Storage Manager storage. The HSM 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.

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 low threshold option, take into consideration that the .SpaceMan directory (which contains system-specific control files created by Tivoli Storage Manager) occupies up to 10% of file system space. Files from this directory are not eligible for migration. You receive a ANS9094W message during automigration when you set the low threshold to a value that is less than 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 dsmscout 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.

Premigration percentage

The premigration percentage controls premigration of additional files after threshold or demand migration completes. The names of any migrated files are removed from the current migration candidates list.

To premigrate the next files that are listed in the migration candidates list, copies of the files are sent to Tivoli Storage Manager storage, and the original files remain as premigrated files on your local file system.

The default for the premigration percentage is the difference between the percentages that you specify for the low and high thresholds. The default is not greater than the low threshold. The default percentage premigrates enough files to make the next occurrence of threshold migration faster. For example, if the high threshold is 90 percent and the low threshold is 80 percent, the premigration percentage is 10 percent. When space usage drops to 80 percent, additional files premigrate until at least 10 percent of the occupied space on your file system contains premigrated files that are listed at the beginning of the current migration candidates list. The next time threshold migration is required, the HSM client replaces those files with stub files on your local file system. This quickly reduces space usage to a low threshold without requiring additional time to copy the files to Tivoli Storage Manager storage.

If demand migration is required, having your eligible files already premigrated hastens that process as well. The HSM client quickly releases at least ten percent of the space on your local file system and migrates any additional files that are necessary to return space usage to the low threshold.

During premigration, the HSM client skips any files that are premigrated and premigrates only those files that are required to reach the premigration percentage. Increase or decrease the premigration percentage when you want to change that percentage.

Premigration percentage can be configured manually. As a result, premigration percentage does not adapt to any low or high threshold modifications in the way

the default premigration percentage value does. 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 HSM client does not premigrate any files after threshold or demand migration completes.

If the premigration percentage equals the percentage that you specify for the low threshold, the HSM client premigrates all remaining files in your file system that are currently eligible for automatic migration.

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 HSM 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 Tivoli Storage Manager administrator to determine whether there are any restrictions on the amount of data that you can migrate and premigrate to storage.

Stub file size

A stub file contains information that is necessary to locate and recall a migrated file. It can contain leading bytes of data called *leader data* from your original file. If you access only this data but you do not modify it, the migrated file is not recalled from Tivoli Storage Manager storage. Storing leader data in stub files is especially useful if you frequently run one or more programs that read only the information located at the beginning of a large number of files.

When selecting a stub file size, consider the advantage of storing additional leader data (files are not recalled if only leader data is accessed and not modified), over the disadvantage of using additional space on your local file system for stub files (larger stub files consume more storage space on your local file system).

For all file system types on all platforms, the maximum value for a stub file size is one gigabyte.

For HSM clients on AIX JFS2 file systems, the minimum stub file size is 511 bytes. You can select a stub file size that is smaller than your file system fragment size.

However, this is an inefficient use of space on your file system. Although the HSM client uses only the number of bytes that you specify for stub files, your operating system allocates one fragment for each stub file. The excess number of bytes remain unused. The default stub size is 4095.

For HSM clients on AIX GPFS and Linux x86/x86_64 GPFS file systems, valid stub file sizes are 0 or multiples of the file system fragment size. The default stub size is 0.

For HSM clients on Solaris VxFS and HP-UX VxFS file systems, valid stub file sizes are a multiple of the file system blocksize. The default stub size is the file system block size.

Minimum stream file size

In an AIX GPFS, AIX JFS2, Linux x86/x86_64 GPFS, and Solaris VxFS file systems environment, 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 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. For additional information also refer to “Streaming recall mode” on page 77.

Minimum partial file size

In an AIX GPFS or Linux x86/x86_64 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.

For additional information also refer to “Recalling migrated files overview” on page 3 and Chapter 5, “Migrating files,” on page 59.

Deactivating space management

Use a command or the HSM GUI to deactivate space management on a single file system or globally on all space managed file systems. Deactivate space management prior to before you or your Tivoli Storage Manager 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 first.
- You can access only resident and premigrated files on a deactivated file system.
- You can also perform file system actions by using the **Actions** → **File System** menu or the menu that opens when you right-click a file system in the file system table.
- If your administrator exports migrated files from one server and imports them to another, update the `dsm.sys` file so that the client node contacts 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 prior to performing the import and export tasks.

Use one of the following commands to deactivate space management on a single file system or globally across all space managed file systems (See the **dsmmigfs** command information, “dsmmigfs deactivate, reactivate, and remove” on page 116):

```
dsmmigfs deactivate /home
dsmmigfs globaldeactivate
```

You can also follow these steps to deactivate space management from the HSM GUI:

1. From the Manage Resources window, click a client node in the **Client Nodes** tree. If the client node is not connected, then you need to sign on to the client node. The list of file systems on a client node is displayed in a table in the work area next to the **Client Nodes** tree.
2. Select an Active file system from the work area, and click the **Deactivate** button. The status of the file system becomes **Deactivated**.

Reactivating space management

If you deactivated space management from your file systems, you can reactivate space management at any time using the command line or the HSM GUI. If you globally deactivated your file systems, you can also globally reactivate them.

The state of the file system becomes active after you reactivate it. Files in the active file system can again be migrated to and recalled from the Tivoli Storage Manager server.

You can reactivate a single file system or globally reactivate deactivated file systems with the commands:

```
dsmmigfs reactivate /home
dsmmigfs globalreactivate
```

See “dsmmigfs deactivate, reactivate, and remove” on page 116.

To reactivate a deactivated file system using the HSM GUI. The state of a client node is displayed in the **Status** column of the table in the client nodes work area in the Manage Resources window. The state of a client node can be one of the following:

- Not connected
- Active
- Global deactivated
- Not manageable
- Not managed

You can also perform file system actions by using the **Actions-> → File System** menu or the menu that opens when you right click a file system in the file system table.

1. From the Manage Resources window, right-click a client node in the **Client Nodes** tree. If the client node is not connected, then you need to sign on to the client node. The list of file systems on a client node appears in a table in the work area next to the **Client Nodes** tree.
2. If the node was globally deactivated, click **Global Reactivate** from the pop-up menu. Otherwise, to activate a deactivated file system from the work area, select the file system, and click the **Reactivate** button.

Removing space management

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

When you remove space management from your file system, the HSM client performs the following:

- 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. For more information, see “Resolving orphaned stub files” on page 81.
- Determines how much space is required to recall all migrated files. If there is not enough space, the HSM 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 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.

Removing space management using the `dsmmigfs` command

To remove the space management from the `/home` file system, issue the command:

```
dsmmigfs remove /home
```

See “`dsmmigfs` deactivate, reactivate, and remove” on page 116.

Removing space management using the HSM GUI

To remove space management from an active file system using the HSM GUI, follow these steps:

1. From the Manage Resources window, click a client node in the **Client Nodes** tree. If the client node is not connected, then you need to sign on to the client node. The list of file systems on a client node appears in a table in the work area next to the **Client Nodes** tree.
2. Select an **Active** file system from the table in the work area, and click the **Unmanage** button. If migrated data is on the server, a warning dialog window opens to inform you that all migrated files will be recalled from the server.

Note: You can also perform file system actions by using the **Actions File System** menu or the menu that opens when you right-click a file system in the file system table.

Chapter 5. Migrating files

The HSM client automatically migrates files to Tivoli Storage Manager storage when space is needed on your local file system. On your local system, stub files replace these migrated files. Your system administrator can determine the size of these stubs on a per file system basis.

At any time, you can migrate particular files. For example, use selective migration if you run threshold migration and some files do not migrate to storage because their last access dates make them ineligible for migration. The files you specified then migrate immediately to the Tivoli Storage Manager server storage.

Note: Any application that touches a file causes that file's last access date to change. This change influences the file processing of Hierarchical Storage Management, since the last access date is one of the factors that determines when a file becomes eligible for migration.

You can use the Tivoli Storage Manager *serveLastAccessDate* option during a backup or archive operation to specify whether to reset the last access date of any specified files to their original value following the backup or archive operation. By default, the IBM Tivoli Storage Manager client *does not* reset the last access date of any backed up or archived files to their original value following the backup or archive operation.

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

- It is a regular file that you previously have not migrated. Character special files, block special files, FIFO special files (named pipe files), or directories are not migrated.
- It is a resident or premigrated file on a file system for which space management is active.
- It is not excluded from space management in your include-exclude options file.
- Its file size is greater than both the stub file size plus one byte and the file system block size.
- It meets management class criteria.

You can migrate any file in a set of hardlinked files that you did not exclude from space management and to which you assigned a management class permitting automatic or selective migration.

For AIX GPFS and Linux x86/x86_64 GPFS file systems only:

Newly-created files that you migrate either automatically or selectively must be older than two minutes before you can migrate them. Migrating newly-created files less than five minutes old might display incorrect results (resident size) when you use the **dsmdf** and **dsmdu** commands because the GPFS is not synchronized on all nodes when you migrate files. The **dsmdf** command will display correct results after GPFS synchronization and after the next reconciliation of the file system.

Automatic file migration

To ensure that free space is available on your local file systems, the HSM 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 HSM 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 storage. When space usage reaches the low threshold that you set for a file system, migration stops. For example, if you set the high threshold for your file system to 80 percent and the low threshold to 70 percent, files begin migrating to 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. Additional files will be premigrated as specified with the option *pmpercentage*.

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 ten 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, if you attempt to copy a very 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 storage. As space becomes available, the process continues to copy the file to your file system (with the exception of AIX GPFS and Linux x86/x86_64 GPFS which stops the process). You only get an out-of-space error message on AIX GPFS and Linux x86/x86_64 GPFS.

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

There is a potential impact on applications which depend on a timely response to write requests. The time until the user's process continues depends on how fast objects are migrated from the file system to create free space and on the configuration of the HSM client. For instance, if the HSM client is configured to require a backup before migration, and not enough objects are backed up, the migration process needs to wait until enough objects are backed up. In order to avoid a long suspension of the process that initiated demand migration, you must have enough eligible candidates.

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 and prioritizes the files based on the file age, the number of days since a file was last accessed, and on the file size.

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

- Reside 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 on every file. The information is used to generate a list of files that is used for threshold and demand migration. Prior to migration, each file is checked to determine if it is still eligible for migration.
- The **`dsmrecall`** and **`dsmmigrate`** commands update the CFI. When a file is restored as a migrated file, the file information is stored in the CFI.

Commands and options that can be used to modify or monitor automatic migration are:

- The *`minmigfilesize`* option in the `dsm.sys` file to specify the minimum file size for a file to be eligible for automatic migration.
- The **`dsmscoutd scanplan`** command to 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 to change the number of scans that can run in parallel in the daemon.
- The *`maxcandidates`* number to improve the performance of automatic migration by lowering the number to 100.

Important: Do not increase the number of *`maxcandidates`* higher than 1000. This will slow down the automatic migration. For best results on your system, select numbers from 10 to 500. The optimal number to use depends on the performance of the file system and operating system.

File premigration

For faster migration, the HSM client prepares files for automatic migration using a process called premigration. Premigrated files are copied to storage beyond the low threshold that you set, but the original files remain intact on your local file system. When free space is again needed on your local file system, premigrated files become migrated files.

The HSM client verifies that files did not change since they became premigrated. When your premigrated files migrate to storage, stub files replace them on your local system.

The HSM client premigrates files each time it completes automatic migration if:

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

The premigration percentage represents the amount of free space on your file system containing premigrated files that are the next eligible candidates for migration. The default for the premigration percentage is the difference between the percentage that you set for the high threshold and the percentage that you set for the low threshold for your file system. You can change the premigration percentage at any time.

Manually starting threshold migration

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

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. If multiple storage pools in a file system reach high threshold, starting automigration on one pool until it reaches low threshold does not result in a low threshold for the 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.

If additional files are eligible for migration and you did not exceed the premigration percentage that you set for your file system, additional files premigrate after the threshold migration process completes. The status of the premigration process displays in the Threshold Migration Status window.

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 command (See “dsmautomig” on page 107 for more information about this command.):

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

Migrating selectively using the `dsmmigrate` command

Use selective migration (the `dsmmigrate` command) to move specific files from your local file systems to 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 storage to free additional space on your local file system.

Note: On large file systems, selective migration can take a while to complete.

The HSM 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 (`atime`) for the file does not change. Unlike automatic migration which uses the age, the number of days since you last accessed a file has no effect on whether your file is eligible for selective migration.

An eligible file must meet the following management class requirements:

- The management class that you assigned to the file permits selective migration.
- A current backup version of the file exists on your migration server if the management class requires one.

Use the following examples to know how to formulate the `dsmmigrate` command in your environment (see “`dsmmigrate`” on page 124):

- To migrate a file named `proj1rpt` from the `/home/proja` directory, issue the command:

```
dsmmigrate /home/proja/proj1rpt
```

- To display information about your migrated files, use the `-detail` parameter with the `dsmmigrate` command.
- To migrate files in any subdirectory below the specified directory that matches the file specification, use the `-recursive` parameter. For example, to migrate all files in a directory named `/migfs2/test/dir1` and in all of its subdirectories, enter:

```
dsmmigrate -R /migfs2/test/dir1
```

- The `dsmmigrate` command can be invoked from a shell application or a script to build a list of files to be recalled (see Chapter 1, “Space management client for UNIX and Linux overview,” on page 1). This file list can then be passed directly to the HSM client. For example, the following command migrates all files owned by user `ibm`:

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

GPFS policy-driven migration

GPFS 3.2 provides a disk based Information Lifecycle Management (ILM) implementation with the storage pool concept. The Tivoli Storage Manager for Space Management 6.1 Client integrates tape into the GPFS ILM solution. With Tivoli Storage Manager 6.1 and GPFS V3.2, the HSM client can be used with the GPFS policy based storage management features to do ILM to tape and external storage pools.

GPFS 3.2 introduces a new feature called external storage pools. You can set up external storage pools and GPFS policies allowing the GPFS policy manager to coordinate file migrations from a native GPFS online pool to external pools on the Tivoli Storage Manager server. The GPFS policy manager invokes the migration through the HSM client command line interface.

To set up the HSM client to run with the GPFS 3.2 policy manager, perform the following steps:

- Install GPFS 3.2 and Tivoli Storage Manager 6.1.
- Enable the GPFS configuration according to the policy-driven demand migration and to high and low threshold migration by setting the GPFS `enableLowspaceEvents` option to *yes*.

To see the current value of this option, issue the command:

```
mmfsconfig | grep enableLowspaceEvents
```

If this value is *no* or not shown, then change this value to *yes* by issuing the command:

```
mmchconfig enableLowspaceEvents=yes
```

If this option is set to *yes*, then demand migration (where high and low thresholds were set with HSM) is disabled. To enable demand migration again, set the GPFS `enableLowspaceEvents` option to *no*.

- Disable the HSM automigration capability with either HSM client CLI command:

```
dsmmigfs add -HT=100 /yourFS  
dsmmigfs update -HT=100 /yourFS
```

For more information on these commands, see “`dsmmigfs add and update`” on page 112.

- Add an external HSM storage pool to the placement policy file.

For example:

```
RULE EXTERNAL POOL  
    'hsmpool1' EXEC '/usr/lpp/mmfs/samples/ilm/mmpolicyExec-hsm.sample' OPTS '-v'
```

Note: Ensure that the script defined for EXTERNAL POOL under EXEC is executable and installed only on the HSM client nodes.

- Add a threshold migration rule to the placement policy file.

For example:

```
RULE 'HsmData' MIGRATE FROM POOL 'StoragePool1'  
THRESHOLD(90,80,70)  
WEIGHT( CURRENT_TIMESTAMP - ACCESS_TIME )  
TO POOL 'hsmpool1'  
WHERE FILE_SIZE > 1024
```

This example rule creates a rule named `HsmData`. Files are migrated from the online storage pool named `StoragePool1` to the external storage pool named `hsmpool`. The files are chosen based on their last access time and for those larger than 1024 bytes. The migration of files starts when the 90% high threshold is reached. Migration continues until the 80% low threshold is reached. The difference between the low and premigration threshold defines the percentage of data which is premigrated. (80% - 70% = 10 %)

Note: For more details about storage pools, policy settings, and rules read the *GPFS 3.2 Advanced Administration Guide* (SC23-5182-01).

- Install the placement policy with the GPFS CLI command:

```
mmchpolicy /yourFS yourPolicyFileName
```

Note: For more details about the `mmchpolicy` command, see the *GPFS 3.2 Administration and Programming Reference* (SC23-2221-01).

When the settings shown above are applied, GPFS performs the role of the HSM automigration function. The GPFS policy manager starts when the predefined high

threshold for a native GPFS storage pool is reached. The policy manager generates a list of migration candidates (files) and invokes the HSM **dsmmigrate** command. The candidate files in the list are migrated to the external storage pools in the Tivoli Storage Manager server.

Chapter 6. Archive or retrieve and backup or restore overview

You need to understand the specific terminology and processes involved with backing up and restoring files in contrast to archiving and retrieving them and migrating and recalling them. In most cases, you use the backup-archive client for backup, restore, archive and retrieve operations and use the HSM client for migrating and recalling files.

When you backup a file using the backup-archive client, a copy of the file is created on the Tivoli Storage Manager server and the original file remains in your local file system. To obtain a backed file from Tivoli Storage Manager 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 Tivoli Storage Manager storage, that file is removed from your local file system, and if needed, you *retrieve* it from Tivoli Storage Manager storage.

When you migrate file, you move the file to Tivoli Storage Manager 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.

You need to regularly use a Tivoli Storage Manager backup-archive client to guard against loss or corruption of your data, regardless of whether the files are resident, migrated, or premigrated. See IBM Tivoli Storage Manager for UNIX and Linux Backup-Archive Clients Installation and User's Guide for information about backing up and restoring files.

You can back up and migrate your files to the same IBM Tivoli Storage Manager server or to different IBM Tivoli Storage Manager 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 Tivoli Storage Manager server.

To restore stub files rather than backup versions of your files, for example if one or more of your local file systems is damaged or lost, use the backup-archive client **restore** command with the *restoremigstate* option. To restore the stubs of HSM managed files with the backup-archive client, the *dsmrecalld* daemon must be running. Your migrated and premigrated files remain intact on the Tivoli Storage Manager server, and you need only restore the stub files on your local system. However you cannot use the backup-archive client to restore stub files for your migrated files, if they have been backed up before the migration. Instead use the HSM client **dsmmigundelete** command to recreate stub files for any migrated or premigrated files that are lost.

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 this by limiting the number of tape drives which can be used simultaneously by backup and archive operations. Specify a number for the *mountlimit* which is less than the total number of drives available in the library (see *mountlimit* option of the **define devclass** command in the IBM Tivoli Storage Manager Administrator's Reference for your operating system). 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.

If you back up files to one server and migrate them to a *different* server or if you are using different server stanzas for backup and migration, the HSM client cannot verify that current backup versions of your files exist before you migrate them. Use the backup-archive client to restore the actual backup versions only.

Archiving and retrieving files using the backup-archive client

You can archive your files at any time and retrieve them to your local file systems when you need them. Use the Tivoli Storage Manager backup-archive client to archive and retrieve copies of your migrated or premigrated files in the same manner as you would archive and retrieve copies of files that reside on your local file systems.

See IBM Tivoli Storage Manager for UNIX and Linux Backup-Archive Clients Installation and User's Guide for more information about archiving and retrieving files.

Archiving migrated or premigrated files

If you archive a copy of a migrated file to the same server to which it was migrated, the file is copied from the migration destination to the archive destination. Archiving a migrated copy does not recall the file to your local file system. If you archive a copy of a premigrated file to the same server to which it was migrated or to a different server, a copy of the file is sent from your local file system to storage.

If you erase a file from your local file system after you archive a copy of the file, the stub file is deleted from your local file system. However, the migrated copy of the file remains in storage until it expires. The number of days that you specify on the *migfileexpiration* option in your *dsm.sys* file determines file expiration.

If you do not erase the file from your local file system after you archive a copy, the file remains migrated.

Note: For Solaris VxFS, HP-UX VxFS, AIX GPFS, and Linux x86/x86_64 GPFS file systems, if you archive a copy of a migrated file to a server other than the one to which it migrated, the backup-archive client accesses and recalls the file. It resides on your local file system in a premigrated state until it is migrated again, or until it receives resident status.

Retrieving archived files

When you need to retrieve an archived file, you retrieve a copy of a migrated or premigrated file to your local file system by changing the *restormigstate* option in your *dsm.opt* file. The file is restored to your local file system as a resident file.

Note: For AIX GPFS, Linux x86/x86_64 GPFS, Solaris VxFS, and HP-UX VxFS file systems files with ACLs are restored to a resident state, even when you specify *yes* on the *restoremigstate* option in your *dsm.opt* file.

- To recall archived files and remove the file from storage when it expires, set the *restoremigstate* option to *no* in your *dsm.opt* file.

- 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* in your *dsm.opt* file.

Backing up files before or after migration

You need to back up your migrated files to guard against data loss. Use a management class to specify if your files should be backed up before or after migrating them.

If you back up and migrate files to the *same* server, you can assign a management class to files and specify that current backup versions of your files must exist on the migration server before the files migrate. The default management class includes this requirement. The HSM client checks for backup versions of files only on the server to which it migrates your files and if a current backup version of a file does not exist on that server, the file is not migrated.

If you back up files to one server and migrate them to a *different* server, the HSM client cannot verify that current backup versions of your files exist before it migrates them. Any management class that you assign to files must specify that current backup versions are not required prior to 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 and restoring premigrated files

You can back up a premigrated file in the same manner that you back up a resident file. Whether you back up a file to the same server to which it was premigrated or to a different server, a copy of the file is sent from your local file system to storage. The HSM client does not copy the file from the migration destination to the backup destination.

If you back up and migrate files to the *same* server, you can:

- Restore backup versions of premigrated files
- Restore stub files for premigrated files
- Create stub files for premigrated files

If you back up and migrate files to a *different* server, you can:

- Restore backup versions of premigrated files
- Create stub files for premigrated files

For more information about restoring backup versions and stub files for migrated or premigrated files, see *IBM Tivoli Storage Manager for UNIX and Linux Backup-Archive Clients Installation and User's Guide*.

Restoring files

When you restore a file, you copy a version of a backed up file from storage. For most restore operations, you use the backup-archive client. You can use the HSM client **dsmmigundelete** command to restore stub files.

Read through the following steps before taking any action so you fully understand your range of options. For information on restoring just stub files, see (“Recreating stub files using the dsmmigundelete command”).

Each step below shows you a different method for restoring files.

1. If you backup and migrate files to the *same* server, you can restore a backup version if the stub file is lost or corrupted and you want to restore the entire file. Set the *restoremigstate* option to *no* in your *dsm.opt* file. The file becomes a resident file. The migrated copy of your file is removed from IBM Tivoli Storage Manager storage when it expires. Specify the expiration with the *migfileexpiration* option in your *dsm.sys* file.

Note: For more information about these options, see “Restoremigstate” on page 99 and “Migfileexpiration” on page 96.

2. If you backup and migrate files to the *same* server, you can restore a stub file if the stub file on your local file system is lost or corrupted, and you want the file to remain migrated. Set the *restoremigstate* option to *yes* in your *dsm.opt* file.

Note: For AIX GPFS, Linux x86/x86_64 GPFS, Solaris VxFS, and HP-UX VxFS file systems only: Files with ACLs are restored to a resident state, even when you set *restoremigstate* to *yes*.

Recreating stub files using the dsmmigundelete command

If your stub files are erased or corrupted and you backed up and migrated files to a *different* server, you can use the **dsmmigundelete** command to recreate stub files for all eligible migrated files in the file system that you specify.

Note:

- You cannot use the **dsmmigundelete** command to recreate 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 you back up and migrate files to the *same* server and backup versions of your files exist use the backup-archive client to restore your stub files.
- The **dsmmigundelete** command does not support hardlinked files. If you attempt to recreate a stub file for a hardlinked file, a stub file is not recreated *unless* all of the files that are hardlinked together are deleted from your local file system. When one file in a set of hardlinked files is migrated, all of the hardlinked files in the set become stub files. When the **dsmmigundelete** command recreates a stub file for a hardlinked 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 hardlinked set of files.
- For more information about using the **dsmmigundelete** command, see “dsmmigundelete” on page 126.

When you issue the **dsmmigundelete** command:

- The **dsmmigundelete** command creates a stub file containing the necessary information to recall the corresponding file from storage. It 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 recreated stub file. The recall mode for the file is set to normal.
- The HSM client does not create a stub file if a directory path does not exist in your local file system for a migrated or premigrated file.
- A stub file is recreated 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, a stub file is created again for the file with the most recent modification time (mtime).
- If the name of a migrated or premigrated file in storage is the same as the name of a file that currently resides on your local file system, the HSM client replaces the file on your local file system with a stub file only if the modification time for the migrated or premigrated file is newer than the modification time for the file on your local file system.
- The HSM client can recreate a stub file for a migrated file or create a stub file for a premigrated file even if:
 - The file was never backed up.
 - The migrated or premigrated file resides on a different server other than the server on which backup copies of the file reside.
 - 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 HSM client performs the following action for the file system that you specify:

Table 19. *dsmmigundelete* command actions without the *expiring* option

File	Description
Migrated files	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. The migrated file will be unexpired at the server after the dsmmigundelete command is finished.
Pre migrated 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 will be unexpired at the server after the dsmmigundelete command is finished.

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

Table 20. *dsmmigundelete* command actions with the *expiring* option

File	Description
Migrated files	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.
Pre migrated files	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.

Restoring file systems overview

If you lose an entire file system and you attempt to restore backup versions of all your files, including those that are migrated and premigrated, your file system might run out of space. If your file system runs out of space during the restore process, the HSM client must begin migrating files to storage to make room for additional restored files, thereby slowing the restore process.

Note:

- When you restore an entire file system, perform one of these tasks *before* you restore backup versions of your resident files.
- You can restore your file system to its state as of the last incremental backup.
- Premigrated files change to a migrated state.

Rather than restoring backup versions of all your files, do one of the following:

- If the backup versions and migrated files reside on the *same* server, use the backup-archive client **dsmc restore** command and set the *restoremigstate* option to *yes* (the default) in your `dsm.sys` file. See “Restoring a file system backed up and migrated to the same server” for full instructions. The HSM client restores backup versions of resident files but restores migrated and premigrated files to stub files.
- If the backup versions and migrated files reside on *different* servers, or if there are no backup versions of migrated and premigrated files, use the HSM client **dsmmigundelete** command to recreate stub files for migrated and premigrated files. See “Restoring a file system backed up and migrated to a different server” on page 73) for full instructions.

Restoring a file system backed up and migrated to the same server

You can restore your file system if you back up and migrate files to the *same* server, and you have backup versions of your migrated and premigrated files.

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.

1. If needed, reinstall the backup-archive client, API, and HSM client.
2. Follow your operating system instructions to establish the file system again.
3. Mount the file system.
4. Add space management to your file system. For information about adding space management, see “Adding space management to file systems” on page 48.
5. 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, issue:

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

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

Restoring a file system backed up and migrated to a different server

If you need to restore an entire file system and you back up files to one server and migrate them to another or if backup versions are not available for migrated and premigrated files you can follow a specific procedure to restore the file system.

Note:

- For more information about the **dsmmigundelete** command, see “Recreating stub files using the dsmmigundelete command” on page 70.
- For more information about the **dsmc restore** command, see the IBM Tivoli Storage Manager for UNIX and Linux Backup-Archive Clients Installation and User’s Guide

Follow these steps in exact order to restore your file system:

1. If needed, reinstall the backup-archive client, API, and HSM client.
2. Follow your operating system instructions to establish the file system again.
3. Mount the file system.
4. Add space management to your file system. For information about adding space management, see “Adding space management to file systems” on page 48.
5. Enter the backup-archive client **dsmc restore** command with the *dirsonly* option to restore the directory structure of your file system. For example, to restore the directory structure for the /home file system, issue:

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

When you use the *dirsonly* option with the **dsmc restore** command, only those backed-up 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.

When you use the *dirsonly* option with the **dsmc restore** command, only those backed up directories for the file systems that you specify are restored.

6. Enter the HSM client **dsmmigundelete** command to recreate stub files for migrated files and to create stub files for premigrated files. For example, if you want to recreate stub files for all migrated files and create stub files for all premigrated files in the /home file system that were not marked for expiration on the server, issue:

```
dsmmigundelete /home
```

Note:

- When restoring an entire file system, do not use the *expiring* option with the **dsmmigundelete** command unless you want to create stub files for all migrated and premigrated files, whether or not the files were marked for expiration. If you use the **expiring** option, you might create stub files for migrated or premigrated files that were intentionally deleted from your local file system prior to the problem that caused you to lose your file system.
 - The HSM client restores hard links during the restore process if the hard links were backed up.
7. 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:

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

Restoring a disk

To restore a disk in the event of disk loss, you restore different aspects of the disk system so you can run a backup-archive client.

The following is a list of system elements you need to restore after a disk loss:

- Operating system
- Communications software
- Tivoli Storage Manager backup-archive client
- HSM client
- File systems

Restoring your operating system and file systems

If you lose the file system that contains the operating system and communications software, you must recover them *before* you can connect to your server. To minimize the impact of such losses, create a set of installation media that will restore your system to a state that permits contact with the server.

Then, if you can run the backup-archive client, you can recover files. If you can run the HSM client, you can recreate stub files for your migrated and premigrated files. If you lose the file system that contains your IBM Tivoli Storage Manager clients, you must reinstall the clients before you can recover your files.

The installation media should contain:

- A working operating system that permits you to perform basic functions.
- A correctly configured communications program that permits you to establish communications with the server. The files that you need depend upon the communications package that you use. Consult your operating system and communications software manuals for help in setting up your installation media.
- A backup-archive client and an HSM client with customized options files. The command line for those clients is sufficient.

Once you restore your operating system and are ready to restore your file system use the procedure in “Restoring a file system backed up and migrated to the same server” on page 72 if you backup files and migrate files to the same server or “Restoring a file system backed up and migrated to a different server” on page 73 if you backup and migrate files to different servers.

Chapter 7. Recalling migrated files

To return a migrated file to your workstation, access the file in the same way as you would access a file that resides on your local file system, such as double-clicking the file. The HSM recall daemon automatically recalls the migrated file from Tivoli Storage Management storage. This process is referred to as transparent recall.

If you want to return specific migrated files to your local file system, use selective recall. When you selectively recall a file, the HSM client stores it to its originating file system. You can also use partial and streaming recall modes.

For AIX JFS2, AIX GPFS, and Linux x86/x86_64 GPFS, and Solaris VxFS file systems, you can set or change the recall mode, for one or more migrated files, with the **dsmatrr** command. Select can select normal, partial recall, or streaming recall mode. For more information about this command, see “*dsmatrr*” on page 105.

Note:

Transparent recall

Transparent recall automatically returns a migrated file to its originating local file system when you access it. Once recalled, the HSM client leaves the copy of the file in storage, but changes it to a premigrated file because an identical copy exists both on your local file system and in storage. If you do not modify the file, it remains premigrated until it once 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 this by limiting the number of tape drives which can be used simultaneously by backup and archive operations. Specify a number for the *mountlimit* which is less than the total number of drives available in the library (see *mountlimit* option of the **define devclass** command in the IBM Tivoli Storage Manager Administrator’s Reference for your operating system). 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.

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

Use selective recall 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 HSM 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
```

See “`dsmrecall`” on page 128 for more information on the **dsmrecall** command.

Normal recall mode

Normal recall mode is the default for all files. In this case files are recalled completely from Tivoli Storage Manager storage and can be accessed after the recall has finished.

Partial file recall mode

In an AIX GPFS or Linux x86/x86_64 GPFS environment, the HSM 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 using compression will always be recalled completely.

If a file qualifies for partial file recall (if it is larger than *minpartialrecallsizesize*) and compression is also 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 HSM 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 *minpartialrecallsizesize* 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
```

See the “dsmrecall” on page 128 command for more information.

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

- Set the *minpartialrecallsizesize* option of the **dsmmigfs** command to the minimum size (in megabytes) that a file must have to qualify for partial file recall. See “dsmmigfs add and update” on page 112 for more information.
- Set the *recallmode* option of the **dsmattr** command to *partialrecall (p)*. 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

For more information, see “dsmattr” on page 105.

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:
 - AIX GPFS, AIX JFS2, Linux x86/x86_64 GPFS, and Solaris VxFS file systems
 - Read-only operations on the file
- Partial file recall mode takes precedence over streaming recall mode. If a file is smaller than its file system’s *MINPartialrecallsizesize* (as configured via the **dsmmigfs add** or **update** commands), or *MINPartialrecallsizesize* is set to 0, normal or streaming recall mode takes precedence.

You can use the following methods to specify which files HSM should recall using streaming recall mode:

- Set the *RECALLmode* option of the **dsmattr** command to *s* (streaming). This 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. When you set the recall mode to streaming recall, this mode remains associated with your files until you:
 - Change the recall mode
 - Modify the file
 - Recall the file selectively
 - Restore the file

See “dsmattr” on page 105 for more information.

- Set the *MINStreamfilesizesize* option of the **dsmmigfs** command to specify a number to enable or disable an asynchronous recall of migrated files. See “dsmmigfs add and update” on page 112 for more information.

How HSM determines which recall mode to use

The HSM determines the recall mode depending on your settings and what happens to a recalled file.

Table 21 displays the recall mode that is used depending on the:

- Recall mode that you set for a migrated file
- Action that the process takes

Table 21. HSM client determining recall modes

If the recall mode for a migrated file is:	If the process:	The recall mode used is:
Normal	Does not modify the file	Normal recall mode. The file becomes premigrated.
Normal	Modifies the file	Normal recall mode. The file becomes resident.
Streaming	Does not modify the file	Streaming recall mode. The file becomes premigrated.
Streaming	Modifies the file	Normal recall mode. The file becomes resident.
Partial	Does not modify the file	Partial recall mode. The file stays migrated.
Partial	Modifies the file	Normal recall mode. The file becomes resident

Setting the recall mode 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.

Note: The `dsmattr` command is valid only for AIX GPFS, AIX JFS2, Linux x86/x86_64 GPFS, and Solaris VxFS file systems.

See “`dsmattr`” on page 105 for more information about this command.

Chapter 8. Reconciling file systems

To keep your local file systems synchronized with the Tivoli Storage Manager server that you contact for space management services, the HSM client automatically reconciles your file systems at intervals that you set. You, as root user, also can start reconciliation manually.

Note: To perform reconciliation tasks, you must have root user authority.

This topic describes the reconciliation options that you set, the reconciliation tasks that the HSM client performs, and the manual performance of these tasks.

Automatic reconciliation

The HSM 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* file. The default is every 24 hours. To specify how many file systems automatically are reconciled at one time, modify the setting on the *maxreconcileproc* option in your *dsm.sys* file. The default is 3 file systems.

To specify how many days a migrated or premigrated file remains in 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* file. The default is 7 days. When the copies expire, they are removed from the server.

For more information about these options, see Chapter 11, “HSM client *dsm.sys* and *dsm.opt* option reference,” on page 89.

Table 22 describes the tasks that automatic reconciliation performs on files and file systems.

Table 22. 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 if:<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 storage if it expired• Updates the status file
Premigrated files	<ul style="list-style-type: none">• Verifies that premigrated files are still valid• Marks a premigrated file for expiration if:<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 storage if it expired.• Updates the status file

Table 22. Automatic reconciliation tasks (continued)

File type	Reconciliation Tasks
Stub files	Records, in the orphan.stubs file, 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. See “Resolving orphaned stub files” on page 81 for more information.
Status file	Updates the following information in the status file: <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>

Manually reconciling file systems

If you recall a migrated file, modify it, and selectively migrate it to storage, two copies of that file reside in storage. The unmodified copy of the file now is obsolete. You can perform manual reconciliation to synchronize your file system and save space by removing obsolete copies from storage.

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

Note: A value of zero 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 reconcile process runs with the *migfileexpiration* option set to 0, the file can not be recreated 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 perform 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:

- Since the scout daemon also aids in the reconciliation process, it must be running. See “dsmreconcile” on page 130 and “The scout daemon” on page 84 for more information.
- 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 to determine if any orphaned stub files were located. If the orphan.stubs file contains file names, see “Resolving orphaned stub files” on page 81.

For command information see “dsmreconcile” on page 130.

Resolving orphaned stub files

An orphaned stub file is a stub file for which a corresponding migrated file in storage is not located. If orphaned stub files exist in your file systems, the HSM 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.

Possible situations in which stub files might become orphaned include the following:

- 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 Tivoli Storage Manager 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 them from storage. In this case, 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.

Chapter 9. Space management daemons

The HSM client uses the space monitor, recall, scout, watch, and root daemons to manage your file systems automatically. They start when you add space management to your file systems and modify different options.

The space monitor, recall, and scout daemons manage storage pools in a GPFS managed file system.

Note: The watch daemon is for AIX GPFS and Linux x86/x86_64 GPFS only.

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 your *dsm.sys* file.

To reconcile your file systems more frequently or less frequently, change the value on the *reconcileinterval* option in your *dsm.sys* file. See “Reconcileinterval” on page 98 for more information about that option.

When you change the option values that the space monitor daemon uses, the new values are not effective until you restart your system or stop and restart the space monitor daemon.

For AIX file systems only:

The space monitor daemon starts automatically when you mount your file system and add space management to it. If the space monitor daemon stops running, issue the **dsmmonitord** command to start it.

The recall daemon

The recall daemon recalls migrated files from storage to your local file system. The “master” recall daemon starts two more instances by default that do not perform recalls: “distributor” and “receiver”. These instances are used for communication purposes only.

Use *minrecalldaemons* and *maxrecalldaemons* options in your *dsm.sys* file to adjust the number of child recall daemons that you can run at one time to perform recalls for your client node.

If a recall daemon is not running, enter the **dsmrecalld** command to start one. For AIX GPFS and Linux x86/x86_64 GPFS see “dsmwatchd” on page 134.

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 your *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. See “Maxrecalldaemons” on page 94 and “Minrecalldaemons” on page 97 for more information about these options.

The maximum number of recalls that you can set depends on the number of concurrent recalls that normally occur on your system.

If all child recall daemons are busy (including on the source nodes participating in distributed HSM on HSM AIX GPFS and Linux x86/x86_64), another file cannot be recalled until a child recall daemon is available. If a frequently-used application opens several files at the same time, and that application uses all available child recall daemons because all files are migrated, increase the value that you set on the *maxrecalldaemons* option. The accessing application is suspended until a child recall daemon is available.

When you change the option values that the recall daemons use, the new values are not effective until you restart your system, or you stop and restart the recall daemons.

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 migrate, 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.

The watch daemon

If failover is active on your AIX GPFS and Linux x86/x86_64 GPFS node, the watch daemon checks the status of the recall, the monitor, the scout and the root daemon. If any of these daemons end or become corrupted, the watch daemon automatically recovers the failed daemon.

If two or more nodes within a GPFS cluster participate actively in a failover environment, the watch daemon actively takes over the file systems of a failed HSM node if the HSM client can no longer perform its operations locally.

The responsiveness service is part of the watch daemon and it monitors the node responses for failover detection. When a node failure is detected, the service initiates the process of failover.

Nodes are made known to the responsiveness service through the */etc/adsm/SpaceMan/config/DSMNodeSet* file. Each node in this file is in the cluster and is added to the responsiveness service for monitoring.

The responsiveness service is for AIX GPFS and Linux x86/x86_64 GPFS only. It replaces the RSCT Group Services functions. Therefore, creating an RSCT peer domain is no longer necessary.

You cannot start the watch daemon manually. It is started from the `/etc/inittab` file.

The root daemon

The root daemon provides non-root user support for HSM and thus allows non-root users to use several HSM commands. If a root daemon is not running, enter the **dsmrootd** command to start one.

Stopping the space management daemons

Do not use **kill -9** for stopping any space management daemon. All daemons have their own clean up procedure, which will be interrupted in this case. Using **kill -9** will lead to unpredictable and unintentional results. Instead use a series of steps to stop the daemons.

Follow these steps to stop the space monitor, master recall, subordinate recall, or scout daemon:

Note: Stopping a master recall daemon stops all subordinate daemons.

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. For more information, see “**dsmq**” on page 127.
2. Issue the **dsmrm** command to remove each recall process from the queue. For more information, see “**dsmrm**” on page 131.
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. For AIX and Linux x86/x86_64 GPFS file systems, issue the command **dmkilld** to stop the recall daemons. For more information see “**dmkilld**” on page 105.
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. For more information, see “**dsmmigfs stop, start, and restart**” on page 120.

Chapter 10. Scheduling services

Your Tivoli Storage Manager administrator defines a schedule on the server and associates your client node with that schedule to perform backup, archive, or space management tasks automatically at specific times. Central scheduling requires a cooperative effort between a Tivoli Storage Manager server and your client node.

Your administrator also sets server parameters to:

- 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 when it is time to perform scheduled services
- Control how often your client node contacts the server for scheduled work

Scheduling options

Before scheduled services can be performed, 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. For information about setting these options in your `dsm.sys` file, see *IBM Tivoli Storage Manager for UNIX and Linux Backup-Archive Clients Installation and User's Guide*.

Starting the client scheduler

Before scheduled services can be performed, 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. Issue the Tivoli Storage Manager backup-archive client **schedule** command to start the client scheduler.

Note:

- For information about setting these options in your `dsm.sys` file, see *IBM Tivoli Storage Manager for UNIX and Linux Backup-Archive Clients Installation and User's Guide*.
- For more information about starting the client scheduler, see *IBM Tivoli Storage Manager for UNIX and Linux Backup-Archive Clients Installation and User's Guide*.

You can start the client scheduler at any time. The client scheduler runs continuously until you stop the process or log off from your system.

You can also 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 HSM client does not automatically reconcile file systems at specific intervals.

Displaying scheduled services information

You can display information about scheduled services as well as information about completed services.

To display information about scheduled services for your client node, issue the **dsmc query schedule** command. This command is provided with the Tivoli Storage Manager backup-archive client (see IBM Tivoli Storage Manager for UNIX and Linux Backup-Archive Clients Installation and User's Guide for more information on this command).

When you run the **dsmc query schedule** command in the foreground, output from scheduled commands displays on your screen. The output is also directed to the log file `dsmsched.log` in the current directory unless you change the path and file name with the *schedlogname* option in your `dsm.sys` file. When you run the **dsmc query schedule** command in the background, output is directed only to the `dsmsched.log` file. Check this log file to verify that all work completed successfully.

Chapter 11. HSM client `dsm.sys` and `dsm.opt` option reference

The HSM 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.

Note:

- See “Editing `dsm.sys` options” on page 29 for details, considerations, an example `dsm.sys` server stanza, and a summary option table about the `dsm.sys` options.
- See “Editing `dsm.opt` options” on page 32 for details, considerations, and descriptions of the two `dsm.opt` options.

Candidatesinterval

The `candidatesinterval` option specifies how often the `dsmscoutd` daemon searches the file systems for files which you added for space management.

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.

Place this option *at the beginning* of your `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.

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.

Place this option *at the beginning* of your `dsm.sys` file *before* any server stanzas.

Syntax

►►—CHECKFororphans—

No
Yes

—►►

Parameters

No The **dsmreconcile** command checks for orphans only when you remove the HSM 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*. For more information, see “The scout daemon” on page 84.

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.

Place this option *at the beginning* of your `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.

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 Tivoli Storage Manager throughput.

Place this option in the client system options file (*dsm.sys*) *within* a server stanza.

Note: 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.

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.

Place this option *at the beginning* of your *dsm.sys* file *before* any server stanzas.

Syntax



Parameters

servername

Use the *defaultserver* option to specify the name of the Tivoli Storage Manager server to contact for backup-archive services if more than one server is defined in the *dsm.sys* file. By default, Tivoli Storage Manager will contact the server 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 *dsm.sys* overrides *defaultserver* in *dsm.opt*.

Place this option at the beginning of your *dsm.sys* file before any server stanzas.

Errorprog

The *errorprog* option specifies a program to which you want to send a message if a severe error occurs during space management processing.

Do not place this option at the beginning of your *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. For example:

```
errorprog /usr/bin/echo
```

Filelist

Use the *filelist* option with the *dsmmigrate*, *dsmrecall*, *dsmls*, or *dsmattr* shell commands to specify an existing file that contains a list of files to be processed by the shell command you issue. Files are processed in the order they appear in *filelist* file.

Do not place this option at the beginning of your *dsm.sys* file before any server stanzas.

The files listed in the file list must adhere to the following rules:

- Each entry must be a fully or partially qualified path to a file or a relative path
- Each entry must be on a new line
- No entry contains wildcard characters
- Each entry results in the processing of only one file object
- The file name is enclosed in quotes if the file name contains any spaces

- The HSM client ignores any entry that is not valid
- No directories are included in the file lists

The following is an example of a list of acceptable files within a filelist:

```
/home/dir/file1
/usr/tivoli/file1
/fs1/dir2/file3
"/fs2/my files/file4"
"/fs3/file.txt"
```

If the file name (the *filelistspec*) that you specify with the *filelist* option does not exist, the command fails. Tivoli Storage Manager skips any entries in the file list that are not valid files. The HSM client logs errors and processing continues to the next entry.

This option is only available from the command line.

Syntax

▶▶—FILEList=— *—filelistspec—*▶▶

Parameters

filelistspec

Specifies the location and name of the file that contains the list of files to process with the command. Here are some examples:

```
dsmmigrate -filelist=/home/dsn/filelist
dsmrecall -filelist=/home/dsn/filelist
dsmls -filelist=/home/dsn/filelist
dsmattr -filelist=/home/dsn/filelist
```

Inclxcl

Use the *inclxcl* option to define the filename and path of your include-exclude options file.

Syntax

▶▶—inclxcl— *—filespec—*▶▶

Parameters

filespec

Specifies the path and file name for your include-exclude file. See “Modifying the include-exclude options file” on page 35 and “Creating an include-exclude list” on page 36 for more information.

Maxcandprocs

The *maxcandprocs* option specifies the number of parallel threads in the scout daemons that can scan for file systems.

Place this option *at the beginning* of your *dsm.sys* file *before* any server stanzas.

Syntax

▶▶—MAXCANDProcs— —*max_number_of_scout_daemons*————▶▶

Parameters

max_number_of_scout_daemons

Specifies the maximum number parallel threads in the scout daemons that can scan file systems. The range of values is 2 to 20. The default is 5.

Maxmigrators

The *maxmigrators* option specifies the maximum number of parallel migration sessions per file system to the Tivoli Storage Manager server that the **dsmautomig** command can perform. Ensure that you have sufficient resources on the server for parallel migration to occur.

Place this option *at the beginning* of your *dsm.sys* file *before* any server stanzas.

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 Tivoli Storage Manager server.

Syntax

▶▶—MAXMIGRators— —*number*————▶▶

Parameters

number

Specifies the maximum number of parallel migration sessions that you can set. The range is 1 to 20. The default is 5. If this option is changed from the default, a corresponding increase should be made in the Tivoli Storage Manager 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 perform recalls for your client node. During normal operations, if the number of recall daemons that are running at one time is similar to the maximum number that is permitted, increase this value.

For example, if you use an application that opens several files at one time, and these files are migrated, the application can use all available recall daemons. If another process attempting to access a migrated file cannot start a recall daemon because the maximum number of recall daemons has been reached, the process stops until a recall daemon is available.

Place this option *at the beginning* of your `dsm.sys` file *before* any server stanzas.

Syntax

▶▶—MAXRECA11daemons— —*number*—▶▶

Parameters

number

Specifies the maximum number of recall daemons that you can run at one time to perform recalls for your client node. The range is 2 through 99. The default is 20.

Maxreconcileproc

The *maxreconcileproc* option specifies the maximum number of reconciliation processes that the HSM client can start at one time.

Place this option *at the beginning* of your `dsm.sys` file *before* any server stanzas.

Syntax

▶▶—MAXRECOncileproc— —*maxreconcileproc*—▶▶

Parameters

maxreconcileproc

Specifies the maximum number of reconciliation processes that the HSM client can start at one time. The range of values is 1 through 99. The default is 3.

Maxthresholdproc

The *maxthresholdproc* option specifies the maximum number of threshold migration processes that the HSM 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.

Place this option *at the beginning* of your `dsm.sys` file *before* any server stanzas.

Syntax

▶▶—MAXThresholdproc— —*maxthresholdproc*—▶▶

Parameters

maxthresholdproc

Specifies the maximum number of automatic threshold migration processes that the HSM client can start at one time. The range of values is 1 through 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.

Place this option *at the beginning* of your `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 through 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.

Place this option *at the beginning* of your `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.

Minmigfilesize

The `minmigfilesize` option specifies the minimum file size for a file to be eligible for migration.

This option applies to all of the HSM-managed file systems for which you have not specified a file system specific value for `-minmigfilesize` using the `dsmmigfs add` or `update` commands. See “`dsmmigfs add and update`” on page 112 for details.

Place this option *at the beginning* of your `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.

For AIX GPFS and Linux x86/x86_64 GPFS file systems, if you specify the default, the HSM 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 HSM 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 AIX GPFS and Linux x86/x86_64 GPFS, it 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 you can run at the same time to perform recalls for your client node.

Place this option *at the beginning* of your `dsm.sys` file *before* any server stanzas.

Syntax

►►—MINRecalldaemons— *—number—*—————►►

Parameters

number

Specifies the minimum number of recall daemons that you can run at one time to perform recalls. The range of values is 1 through 99. The default is 3.

Optionformat

The *optionformat* option specifies the format to use when you specify HSM client commands.

Place this option in your `dsm.opt` file.

Syntax

►►—OPTIONFormat——Standard—
—Short——►►

Parameters

STandard

Issue HSM client commands in a format similar to backup-archive client commands. This is the default. For example:

```
dsmmigrate -Recursive -Detail /home/user1/file1
```

SHort

Issue HSM client commands in a format similar to your operating system commands. For example:

```
dsmmigrate -Rv /home/user1/file1  
dsmmigrate -v -R /home/user1/file1
```

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.

Place this option *at the beginning* of your `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* through *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.

Place this option in your `dsm.opt` file.

You can restore or retrieve a stub file for a migrated file only when:

- 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 has not expired, the file is restored or retrieved to a stub file, regardless of whether it is marked for expiration.

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*. This affects files that are restored to a GPFS file system with more storage pools than the default system pool.

Tip: The *restoremigstate* option does not support hardlinked files. If you want to restore or retrieve a stub file for a hardlinked file, delete all files from your local file system that are hardlinked together. When one file in a set of hardlinked files is migrated, all of the hardlinked files in the set become stub files. When you enter the **restore** command with the *restoremigstate* option, and restore a stub file for a hardlinked 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 hardlinked set of files.

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 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.

Chapter 12. HSM client command reference

Before using commands, review the how to enter commands and their options.

This chapter provides information about the HSM commands. For information about client system options, see Chapter 3, “Selecting HSM client options,” on page 27.

When you issue commands and options, follow these rules:

- *Do not* precede HSM commands with **dsmc**. Each HSM client command is a separately-executable command.
- Issue the complete command name in lowercase letters. You *cannot* use uppercase letters or an abbreviation for a command name.
- Issue all commands in a Solaris environment from the Global Zone.
- 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
- Issue 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.

Standard and short option formats

Some commands have options. These options are specified differently dependent upon the value of the optionformat option.

Set the *optionformat* option to standard or short in the client user options file (dsm.opt). The default value is standard.

For optionformat=standard, follow these rules when issuing an option:

- Uppercase letters in each option description indicate the minimum abbreviation that is permitted. 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 -candidatelist -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
```

If optionformat=short, follow these rules when issuing an option:

- Use the short name of the option.
- Issue the short name options in any combination of uppercase and lowercase letters. Short name options are not case-sensitive.

- Precede each short name option with a hyphen (-). For example:
dsmmigquery -c -v /home
- You do not have to separate the short name options. If the options are not separated, precede the initial options with a hyphen (-). For example:
dsmmigquery -cv /home
- Issue one or more options in a command in any order before or after a file, directory, or file system specification.
- If the option defines a value, do not separate the option name from the value. For example:
dsmmigfs update -t90 /home

The table below gives examples of the standard and short name formats.

Table 23. Option format examples: options without values

Option format	Command
optionformat=standard	dsmmigrate -recursive -detail /home/user1/file1 dsmmigrate -rec -det /home/user1/file1 dsmmigrate -r -d /home/user1/file1 dsmmigrate /home/user1/file1 -r -d
optionformat=short	dsmmigrate -R -v /home/user1/file1

The table below gives examples of the standard and short name formats when a value is defined.

Table 24. Option format examples: options with values

Option format	Command
optionformat=standard	dsmmigfs update -hthresold=90 -lthresold=20 /home dsmmigfs update -hthres=90 -lthresold=20 /home dsmmigfs update -ht=90 -l=20 /home dsmmigfs update /home -ht=90 -l=20
optionformat=short	dsmmigfs update -t90 -l20 /home dsmmigfs update -t90l20 /home dsmmigfs update /home -t90l20

Displaying command-line help

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* or *-h* 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 less than 72 characters causes sentences that are 72 characters wide to wrap to the next line. This 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 25. 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 108 for more information about this command.
dsmls	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>dsmls /home/user1/*</pre> See “dsmls” on page 111 for more information about this command.
dsmdu	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: <pre>dsmdu -Allfiles /home/user1/proj1</pre> See “dsmdu” on page 109 for more information about this command.
dsmmigfs	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: <pre>dsmmigfs query /home</pre> See “dsmmigfs add and update” on page 112 for more information about this command.
dsmmigundelete	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: <pre>dsmmigundelete /home</pre> See “dsmmigundelete” on page 126 for more information about this command.

HSM command summary

Table 26 provides an alphabetical list of the HSM client commands, a brief description of each command, and the command page number.

Table 26. HSM command summary

Command and location	Description
dmkilld	Only valid on AIX GPFS and Linux x86/x86_64 GPFS Stops the master recall daemon and all of its children, and interrupts all active recalls. See “dmkilld” on page 105.

Table 26. HSM command summary (continued)

Command and location	Description
dsmattr	Only valid for AIX JFS2, AIX GPFS, Linux x86/x86_64 GPFS, and Solaris VxFS Sets or displays the recall mode for a migrated file. See “dsmattr” on page 105.
dsmautomig	Starts parallel migration sessions for a file system. See “dsmautomig” on page 107.
dsmddf	Displays space usage information for a file system. See “dsmddf” on page 108.
dsmdu	Displays space usage information for files and directories. See “dsmdu” on page 109.
dsmls	Lists files in a directory and displays file state. See “dsmls” on page 111.
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 112.
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 116.
dsmmigfs sdrreset, enablefailover, disablefailover	Manages recovery from partial system failure (GPFS only). See “dsmmigfs sdrreset, enablefailover, disablefailover” on page 118.
dsmmigfs query	Displays current space management settings for a file system. See “dsmmigfs query” on page 117.
dsmmigfs globaldeactivate, globalreactivate	Deactivates or reactivates space management for a HSM-managed client node. See “dsmmigfs globaldeactivate, globalreactivate” on page 119.
dsmmigfs rollback	Only valid for AIX GPFS and Linux x86/x86_64 GPFS 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 120.
dsmmigfs stop, start, restart	Starts or starts the HSM daemons. See “dsmmigfs stop, start, and restart” on page 120.
dsmmigfs takeover	Transfers the HSM management of a file system to a HSM client node within the same local GPFS cluster. See “dsmmigfs takeover” on page 121.
dsmmighelp	Displays online help for commands. See “dsmmighelp” on page 122.
dsmmigquery	Displays space management information. There are many space management and backup-archive client shared options available using the dsmmigquery -o command. See “HSM and backup-archive client dsmmigquery command shared options” on page 123 for a list of these shared options. See “dsmmigquery” on page 122.
dsmmigrate	Moves selected files from your local file system to Tivoli Storage Manager storage. See “dsmmigrate” on page 124.
dsmmigundelete	Recreates deleted stub files. See “dsmmigundelete” on page 126.
dsmmonitor	Starts the space monitor daemon. See “dsmmonitor” on page 127.
dsmq	Displays information, including recall IDs, for all files that are currently queued for recall. See “dsmq” on page 127.
dsmrecall	Moves selected files from storage to your local file system. See “dsmrecall” on page 128.
dsmrecalld	Starts the recall daemon. See “dsmrecalld” on page 130.
dsmreconcile	Synchronizes the client and server. See “dsmreconcile” on page 130.
dsmrm	Removes a recall process from the recall queue. See “dsmrm” on page 131.

Table 26. HSM command summary (continued)

Command and location	Description
dsmrootd	Starts the root daemon. See “dsmrootd” on page 132.
dsm scoutd	Starts, stops and restarts the scout daemon and shows file system information. See “dmscoutd” on page 132.
dsmsetpw	Changes the Tivoli Storage Manager password for your client node. See “dsmsetpw” on page 133.
dsmwatchd	Only valid on AIX GPFS and Linux x86/x86_64 GPFS Manages failover activities for your HSM client node. See “dsmwatchd” on page 134.

dmkilld

The **dmkilld** command stops the master recall daemon and all of its children and interrupts all active recalls.

Valid on AIX GPFS and Linux x86/x86_64 GPFS only

Syntax

▶—dmkilld—▶

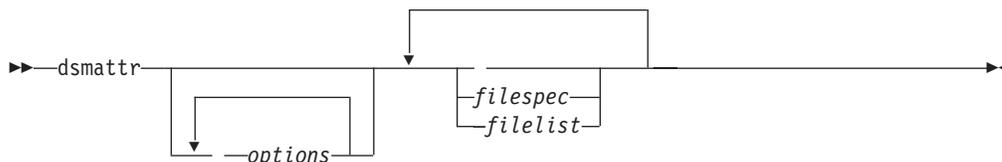
dsmattr

The **dsmattr** command sets or changes the recall mode for one or more migrated files. The recall mode determines how the HSM 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 as long as the file remains migrated.

Valid on AIX JFS2, AIX GPFS, Linux x86/x86_64 GPFS, and Solaris VxFS file systems only

See Chapter 7, “Recalling migrated files,” on page 75 for more information.

Syntax



Parameters

options

—RECALLmode=value

Sets a recall mode for one or more migrated files. If you do not use the **—RECALLmode** option, the HSM client displays the current recall mode for the files that you specify. The values that you can select are:

Value	Description
-	Indicates that the file has not migrated.
<i>Normal (n)</i>	Recalls the migrated file to its originating file system. This is the default. If the file is not modified, it becomes a premigrated file. If the file <i>is</i> modified, it becomes a resident file.
<i>Partialrecall (p)</i> (AIX GPFS, Linux x86/x86_64 GPFS only)	Specifies that the file should be recalled using partial file recall, regardless of its size.
<i>Streaming (s)</i>	Specifies that you want to enable an asynchronous recall of migrated files. The recalled portion of the file can be accessed while the file is recalled. This parameter is valid for read-only operations on the file.

-RECURsive or -R

Sets or displays the recall mode for migrated files in the directory and subdirectories that you specify.

filespec

Important: 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 HSM 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

Processes a list of files.

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>
Enable an asynchronous recall of migrated files in the /home/user2/ directory.	<code>dsmattr -recall=streaming /home/user2/</code>
Display the recall modes that are assigned to all files in the current directory.	<code>dsmattr</code>
Display all files in the filelist named /tmp/filelist.	<code>dsmattr -filelist=/tmp/filelist</code>

dsmautomig

The **dsmautomig** command starts parallel migration sessions to the Tivoli Storage Manager 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 Lotus Notes® Tivoli Storage Manager server.
- If a current backup version exists.

If the LANG environment variable is set to C, POSIX (limiting the valid characters to those with ASCII codes less than 128), or other values with limitations for valid characters, the HSM client skips files which have file names containing invalid characters with ASCII codes higher than 127. 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 HSM 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 (or another SBCS 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 Tivoli Storage Manager server for parallel migration. Do not set the *maxmigrators* option higher than the number of sessions that the Tivoli Storage Manager server can use to store data. Start threshold migration manually to lower space usage on your file system *before* it reaches the high threshold that you set.

Hidden directories and files are included in automatic migration. These can be excluded from automatic migration by adding the hidden directories or files to the exclude list in the dsm.opt file.

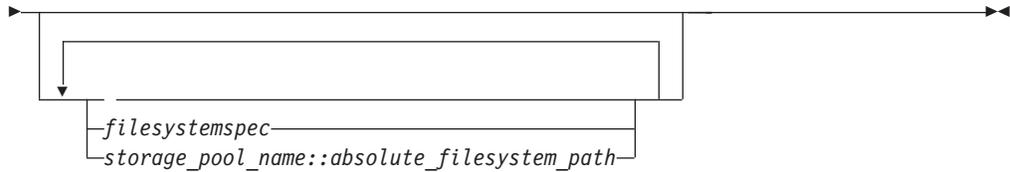
The scout daemon (dsmscoud) should be running if you start the **dsmautomig** command manually. Otherwise, the **dsmautomig** command might not be able to complete the migration if it runs out of candidates from the candidates list.

The **dsmautomig** command must be found with the PATH variable, or the dsmonitord daemon cannot perform threshold migration.

Note: The Tivoli Storage Manager for Space Management client does not migrate contents of symbolic links.

Syntax

►—dsmautomig——detail—→



Parameters

-detail or -v

Displays information about migrated files.

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 located in the *absolute_filesystem_path* that are to be migrated automatically.

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 named <code>silver</code> and <code>gold</code> for the <code>/fs1</code> file systems.	<code>dsmautomig /silver::/fs1 gold::/fs1</code>
Start threshold migration for all storage pools in the <code>/fs2</code> file systems and for the storage pool named <code>gold</code> in the <code>/fs1</code> file systems.	<code>dsmautomig gold::/fs1 /fs2</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 on:

- 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

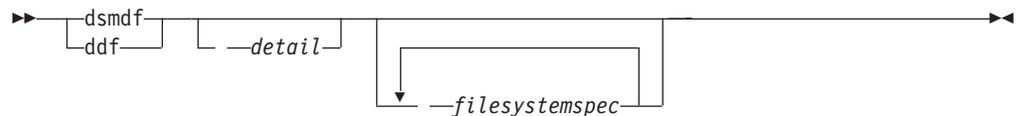
Only migration and recall processes dynamically update status information for your file systems. If any other process changes the state of a file, the change is not reflected in the information that the **dsmdf** command displays until reconciliation is run.

For AIX GPFS and Linux x86/x86_64 GPFS file systems only:

Migrating newly-created files less than five minutes old might display incorrect results (resident size) when you use the **dsmdf** and **dsmdu** commands. 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 a deviation from an assumed disk usage if many small files are migrated and the block size is high.

Note: You can only display information about mounted file systems. If a file system is space managed but not mounted, it will not show up within the output of the command

Syntax



Parameters

detail or -v

Displays information about file systems which each value displayed on its own line. Values representing amount of space will be displayed in kilobytes only.

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>

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 Tivoli Storage Manager 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.

For AIX GPFS and Linux x86/x86_64 GPFS file systems only:

Migrating newly-created files of less than five minutes might display incorrect results (resident size) when you use the **dsmdf** and **dsmdu** commands. 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.

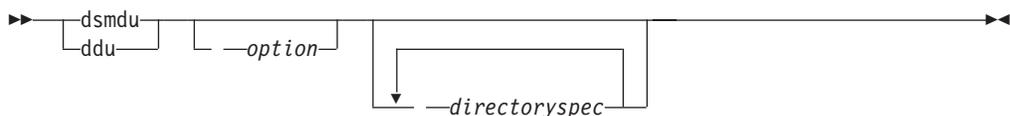
Attention: Running **dsmls** or **dsmdu** on remotely mounted GPFS file systems might show incorrect values. See Table 27 for an example of this situation.

Table 27. HSM management for locally and remotely mounted file systems

Cluster A	Cluster B
HSM client installed	HSM client installed
File system A	File system A (from Cluster A remotely mounted on Cluster B)
Mounted locally	Mounted remotely
HSM-managed by A	Not HSM-managed by B

Only file systems of the local cluster can be managed by HSM. The local Cluster A file system is HSM-managed by Cluster A, but cannot be HSM-managed by Cluster B, although it is remotely mounted on Cluster B. The 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 HSM client displays the number of 1 KB blocks that the specified directory and each of its subdirectories use.

-Allfiles or **-a**

Displays the number of 1 KB blocks that each file in the specified directory and each of its subdirectories use.

-Summary or **-s**

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>

Task	Command
Display space usage information for each file in the /migfs2/test directory and in all of its subdirectories.	<code>dsmdu -a /migfs2/test</code>
Display the total number of 1 KB blocks that the /migfs2/test 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
- 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 be any of the following: 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.

Attention:

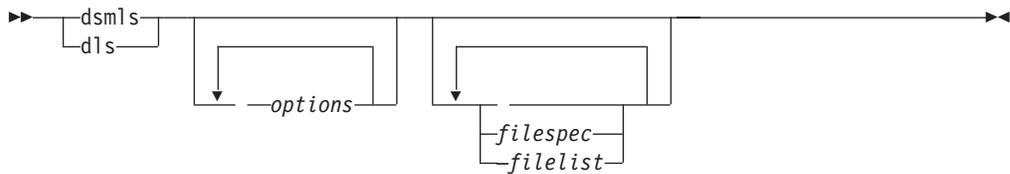
- Running **dsmls** or **dsmdu** on remotely mounted GPFS file systems might show incorrect values. See Table 28 for an example of this situation.

Table 28. HSM management for locally and remotely mounted file systems

Cluster A	Cluster B
HSM client installed	HSM client installed
File system A	File system A (from Cluster A remotely mounted on Cluster B)
Mounted locally	Mounted remotely
HSM-managed by A	Not HSM-managed by B

Only file systems of the local cluster can be managed by HSM. The local Cluster A file system is HSM-managed by Cluster A, but cannot be HSM-managed by Cluster B, although it is remotely mounted on Cluster B. The DMAPI interface does not work for remotely mounted file systems, so incorrect values might be reported by **dsmls** or **dsmdu**.

Syntax



Parameters

options

Use any of the following options:

-Noheader or **-n**

Omits column headings from the output for this command.

-Recursive or **-R**

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

Processes a list of files.

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 <code>tf</code> in the <code>/migfs2/test</code> directory.	<code>dsmls /migfs2/test/tf*</code>
List all files in the filelist named <code>/tmp/filelist</code> .	<code>dsmls -filelist=/tmp/filelist</code>

dsmmigfs add and update

The **dsmmigfs** command adds space management to your file system or updates space management settings for your file system.

You must have root user authority to use this command.

For AIX GPFS and Linux x86/x86_64 GPFS file systems only: Before you run **dsmmigfs add filesystemName**, ensure that file system is mounted and enabled for DMAPI management. Issue the following commands:

For AIX GPFS and Linux x86/x86_64 GPFS:

```
/usr/lpp/mmfs/bin/mm1sfs DevicePath -z
```

If it is required, change the value to:

```
/usr/lpp/mmfs/bin/mmchfs DevicePath -z yes
```

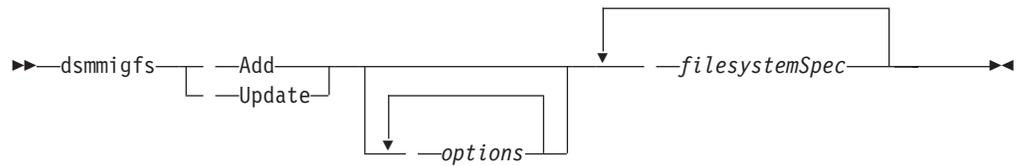
Run only one **dsmmigfs** command within the local GPFS node set at the same time.

You cannot add space management to your root (/), /tmp, /usr, or /var file systems.

JFS2 Encrypted File System is not supported. Adding space management to a JFS2 EFS through the **dsmmigfs add** command displays an error message.

Note: You can only query 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 only for files that are migrated after you make the 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 you make the change.

If you reduce the quota, and the amount of migrated and premigrated data exceeds the new quota, the HSM client does not migrate any additional files until enough files are recalled during automatic recall or selective recall to drop the total number of megabytes for migrated and premigrated files below the new quota.

filesystemSpec

The file system name to perform the specified action. 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. See Table 29 on page 114 for options and information.

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

Table 29. Space management option settings

Option	Short Name	Description
-HThreshold= <i>n</i>	- <i>tn</i>	The high threshold percentage that you set for space usage on your file systems. Specify a value of 0 through 100 percent. The default is 90 percent.
-Lthreshold= <i>n</i>	- <i>ln</i>	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.
-Maxcandidates= <i>n</i>	- <i>mn</i>	The maximum number of migration candidates the dsmscoutd daemon searches for during one scan period. When this number is reached, dsmscoutd stops scanning. The range of values is 9 through 9999999. The default is 10000.
-MAXFiles= <i>n</i>	- <i>maxfn</i>	<p>The maximum number of files in the HSM-managed file system. The range of values is 0 through 999999999999999. The default is 0.</p> <p>This parameter allocates space based on the maximum number of blocks in the file system. If the specified value is dsmmigfs Add -maxFiles=<i>n</i> /hsmmanagedfs1, less than the number of currently used blocks, an error occurs. However, only these items are honored:</p> <ul style="list-style-type: none"> • Values within the range of the current number of files in the file system on the low side • The total likely number of inodes on the high side <p>You can configure the size of the Complete File Index (CFI). This value can be queried or set with a new dsmmigfs command line "-MAXFiles=<i>n</i> (-fn)" parameter. If this parameter is set to 0 (default), the CFI will allocate the maximum required space, equated to the maximum number of blocks possible in the file system. This means that 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 complete file index (CFI) size.</p>
-MINPartialrecallsize= <i>n</i> (AIX GPFS and Linux x86/x86_64 GPFS only)	- <i>nn</i>	Specifies the minimum size (in megabytes) that a file must have to qualify for partial file recall. The range of values is 0 to 999999999. The default value of 0 disables partial file recall for all files.
-MINStreamfilesize= <i>n</i> (Valid for AIX GPFS, AIX JFS2, Linux x86/x86_64 GPFS, and Solaris VxFs)	- <i>zn</i>	Specifies a number to enable or disable an asynchronous recall of migrated files. This number is the number of bytes that must be recalled before HSM starts streaming data to the requesting application (to ensure a steady stream of data). The recalled portion of the file can be accessed while the file is recalled. The range of values is 0 through 999999999. The default is 0.
-Pmpercentage= <i>n</i>	- <i>pn</i>	The percentage of file system space that is available to contain premigrated files. The default is the difference between the percentage that you set for the high threshold and the percentage that you set for the low threshold. Specify a value from 0 through 100 percent.

Table 29. Space management option settings (continued)

Option	Short Name	Description
-Quota= <i>n</i>	-q <i>n</i>	The maximum number of megabytes of data that you can migrate and premigrate from your file system to Tivoli Storage Manager 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.
-STubsize= <i>n</i>	-s <i>n</i>	The size of stub files remaining on your local file systems when files migrate to Tivoli Storage Manager storage. <ul style="list-style-type: none"> • For AIX GPFS and Linux x86/x86_64 GPFS file systems: Specify a value of 0 or a multiple of the file system fragment size. The default is 0. • For Solaris VxFS and HP-UX VxFS file systems: Specify a value of 1 or a multiple of the file system block size. The default is the file system block size. • For all file system types on all platforms: The maximum value for a stub file size is 1 gigabyte. You can specify the stub size using a size modifier like 'k' or 'K' (for kilobyte), 'm' or 'M' (for megabyte) and 'g' or 'G' (for gigabyte).
-MINMigfilesize= <i>n</i>	-y <i>n</i>	Specifies the minimum size (in bytes) that a file must have to qualify for migration. The range of values is 0 through 2147483647. The default is 0. If you specify the default, the HSM 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 non-zero value, the value must be greater than the current file system stub size. Otherwise, the value is ignored. A valid value takes precedence over the global <i>MINMigfilesize</i> option setting specified in the <i>dsm.sys</i> file. For details, please see "Minmigfilesize" on page 97.
-SErver=<servername>	-S <i>server name</i>	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 <i>dsm.sys</i> file. If you do not specify a server name, the default migration server is used. Use a dash (-) to set the server to the default migration server.

Examples

Task	Command
Add space management to the /hsmmanagedfs1 file system. Set the space to the maximum number of blocks in the file system.	<code>dsmmigfs Add -MAXFiles=0 /hsmmanagedfs1</code>
Update the space management settings for the /hsmmanagedfs2 file system. Set the space (number of blocks) to the specified number.	<code>dsmmigfs update -MAXFiles=10000000 /hsmmanagedfs2</code>
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>
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 256 k (kilobytes).	<code>dsmmigfs Add -HT=80 -L=70 -ST=256k /home</code>

Task	Command
Add space management to more than one file system and accept the default values for all space management settings.	<code>dsmmigfs Add /home /test1 /proj*</code>
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. 	<code>dsmmigfs Update -HT=80 -L=70 -ST=1m /home</code>
Specify the minimum size of files in the /home/user1 file system that are recalled using partial file recall.	<code>dsmmigfs Update -minp=100 /home/user1</code>
Specify that the minimum size for a file that can be migrated from the /home/user2 file system must be at least 1 megabyte.	<code>dsmmigfs Update -minm=1048576 /home/user2</code>

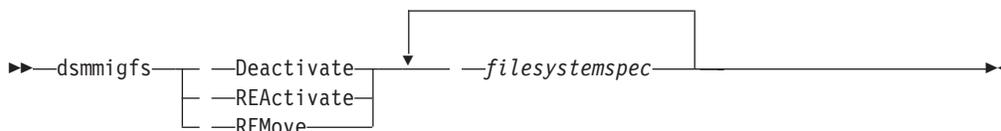
Refer to “Updating settings from the HSM GUI” on page 50 for information about changing your options with the HSM Java GUI.

dsmmigfs deactivate, reactivate, and remove

Use the **dsmmigfs** command with the *deactivate*, *reactivate*, or *remove* parameter to perform those space management actions on a file system.

You must have root user authority to use this command.

Syntax



Parameters

Deactivate

Deactivates space management for a file system. The HSM client cannot perform migration, recall, or reconciliation for 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.

filesystemspec

The file system name that performs the specified action. 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.

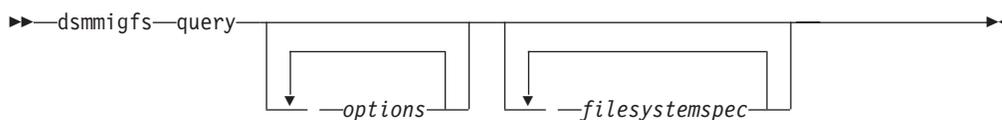
Examples

Task	Command
Deactivate space management for the /home file system.	<code>dsmmigfs Deactivate /home</code>
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 query

The **dsmmigfs** command used with the **query** parameter displays the current space management settings for a file system.

Syntax



Parameters

Query or "q"

Displays the current space management settings for the named file system. The settings include values for the following:

- File system name
- High and low thresholds
- Premigration percentage
- Quota
- Stub file size
- The minimum size (in bytes) that a file must have to qualify for migration
- Server name
- Maximum candidates
- Partial file recall
- Streaming recall
- Server stanza

Note: These settings are only displayed if you specify `-Detail`.

options

-Detail

Displays HSM settings for each file system. Displays options that the HSM client does not display in the standard format, such as the maximum number of candidates that are located during one reconciliation.

For AIX GPFS and Linux x86/x86_64 GPFS file systems only: The **dsmmigfs query** command displays only locally-managed file systems. The **dsmmigfs q -d** command displays information for all space-managed file systems within the GPFS cluster. The HSM client displays the following additional information:

- Node name for each node ID and frame ID for both the owner and the preferred nodes.
- Preferred node. The cluster node where HSM was initially added to the filesystem or the node where the HSM administrator has performed **dsmmigfs takeover**.
- Owner node. The cluster node that is currently managing the filesystem (it can be different from the preferred node after failover).
- Source node. List of nodes that are part of distributed HSM.

The normal **dsmmigfs query** command will not provide any GPFS-specific information (no node set ID).

Note: The **dsmmigfs query -detail** command displays the current status of the local failover environment (either active or inactive).

-failover or -f

Provides a status overview of the failover environment of all HSM-managed cluster nodes. This is valid for AIX GPFS and Linux x86/x86_64 GPFS file systems only. The output displays the status for the node name, and node ID. It can be any of the following:

- Active. The node participates in the failover environment within the local GPFS node set.
- Deactivated by User. You disabled failover using the **dsmmigfs disableFailover** command.
- Deactivated by HSM. The HSM client disabled failover because of an unrecoverable condition.

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. See Chapter 3, “Selecting HSM client options,” on page 27 and “**dsmmigfs sdrreset, enablefailover, disablefailover**” for more information about enabling failover of automigration and recall capabilities to source nodes within a cluster environment.

filesystemspec

The file system name that displays current space management settings. The default is all file systems to which space management has been added.

Examples

Task	Command
Display the current space management settings for the /migfs2 file system.	<code>dsmmigfs query /migfs2</code>

dsmmigfs sdrreset, enablefailover, disablefailover

The **dsmmigfs** command used with the *sdrreset, enablefailover, or disablefailover* parameter manages recovery from partial system failure.

Valid on AIX GPFS and Linux x86/x86_64 GPFS file systems only.

You must have root user authority to use this command.

One HSM client can take over from an HSM client that is involved in a partial system failure if the following conditions are true:

- The failing HSM client node has failover enabled.

- There are one or more additional HSM 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

SDRRReset

Do not use this command during normal operation.

The HSM client attempts to reset potential locking problems in the SDR. If a command or a failover operation ended abnormally, this command will help to achieve a consistent system state.

ENABLEFailover

Activates the node for failover operations within the GPFS cluster.

DISABLEFailover

Deactivates failover operations on the node.

Examples

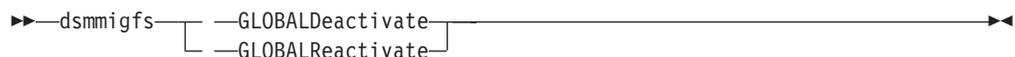
Task	Command
Deactivate failover operations on the HSM client node.	<code>dsmmigfs disableFailover</code>

dsmmigfs globaldeactivate, globalreactivate

Use the `dsmmigfs` command with the `globaldeactivate` or `globalreactivate` parameter to either deactivate or reactivate space management for a HSM-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 HSM client cannot perform migration, recall, or reconciliation for any file system. However, you can update space management settings for file systems, add space management to additional 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.

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 rollback

The **dsmmigfs** command transfers the HSM management of file systems to the preferred node if the node is different from the current owner node.

Valid on AIX GPFS and Linux x86/x86_64 GPFS file systems only.

You must have root user authority to use this command.

Syntax

▶▶—dsmmigfs— —rollback—▶▶

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.

Examples

Task	Command
Transfer to the preferred node.	<code>dsmmigfs rollback</code>

dsmmigfs stop, start, and restart

The **dsmmigfs** with the *stop*, *start*, or *restart* parameter performs the selected command line option on space management daemons.

Valid on AIX GPFS, Linux GPFS systems only

You must have root user authority to use this command.

Use the **dsmmigfs** command with the *start*, *stop*, or *restart* parameter to:

- 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`, which means that `dsm.opt` and `dsm.sys` in the default installation path `/usr/tivoli/tsm/client/ba/bin` will be used.

Syntax

```
▶▶ dsmmigfs —START— —STOP— —RESTART— ▶▶
```

Parameters

START

Starts all HSM daemons on the local client node. The `dsmwatchd` will not be affected.

STOP

Stops all HSM daemons. The `dsmrecall` and `dsmmigrate` processes will also be stopped. The `dsmwatchd` will not be affected.

RESTART

Restarts all HSM daemons. The `dsmrecall` and `dsmmigrate` processes will also be stopped. The `dsmwatchd` will not be affected.

Examples

Task	Command
Start all daemons	<code>dsmmigfs START</code>
Stop all daemons	<code>dsmmigfs STOP</code>
Restart all daemons. For example, let them update the configuration set in your <code>dsm.opt</code> and <code>dsm.sys</code>	<code>dsmmigfs RESTART</code>

dsmmigfs takeover

The `dsmmigfs` command transfers the HSM management of the file system specified by the `fileSpec` parameter to the HSM client node within the same local GPFS node set.

Valid on AIX GPFS and Linux x86/x86_64 GPFS file systems only.

You must have root user authority to use this command.

Syntax

```
▶▶ dsmmigfs —takeover— —filespec— ▶▶
```

Parameters

takeover

The `dsmmigfs` command transfers the HSM management of the specified file system to the HSM 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.

Examples

Task	Command
Transfer the HSM management of the current directory to the HSM client node within the same local GPFS node set.	<code>dsmmigfs takeover /home/filesystem</code>

dsmmighelp

The **dsmmighelp** command displays online help topics from which you can select general help for commands or message information.

Syntax

▶▶ `dsmmighelp` ▶▶

Examples

Task	Command
Display online help for HSM commands.	<code>dsmmighelp</code>

dsmmigquery

The **dsmmigquery** command displays information about migrated files, candidates, and management classes.

You must have root user authority to use this command.

The **dsmmigquery** command displays the following information for one or more file systems:

- Migration candidates list
- Ordered recall list for migrated files
- Available management classes
- Current client and server options
- List of all files in the file system

Output from this command is directed to stdout. Use redirection characters and a file name at the end of the command to redirect the output to a file.

Syntax

▶▶ `dsmmigquery` `-options` `-filesystemspec` ▶▶

Parameters

Options

Select any of these options:

-SORTEDMigrated or -m

Lists all files that you migrated from the file system to Tivoli Storage Manager storage in the most efficient order for recall.

-SORTEDAll or -s

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.

-Mgmtclass or -g

Displays information about each management class that you can assign to your files.

-Detail or -v

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.

-Options or -o

Displays the current settings for your client and server options. This is the default option.

For a list of the **dsmmigquery -o** options that are shared between the HSM client and the backup-archive client, see “Updating settings from the HSM GUI” on page 50.

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 HSM client and the backup-archive client.

The following is a list of the options that are shared between the HSM client and the backup-archive client.

Please see the IBM Tivoli Storage Manager for UNIX and Linux Backup-Archive Clients Installation and User’s Guide for your operating system for more information on using these shared options.

- *asnodename*
- *commmethod*
- *compression*
- *defaultserver*
- *detail*
- *diskbuffsize*
- *enablelanfree*
- *errorlogmax*

- *errorlogname*
- *errorlogretention*
- *exclude*
- *exclude.compression*
- *exclude.file*
- *exclude.file.spacemgmt*
- *exclude.spacemgmt*
- *incl excl*
- *include*
- *include.compression*
- *include.file*
- *lanfreecommmethod*
- *lanfreetcport*
- *lanfreeshmport*
- *mailprog*
- *makesparsefile*
- *nodename*
- *passwordaccess*
- *passworddir*
- *servername*
- *shmport*
- *skipacl*
- *tcpbuffsize*
- *tcpnodelay*
- *tcpport*
- *tcpserveraddress*
- *tcpwindowsize*

dsmmigrate

The **dsmmigrate** command selects specific files from your local file system and migrates them to a Tivoli Storage Manager server.

Attention: On large file systems, selective migration can take a while.

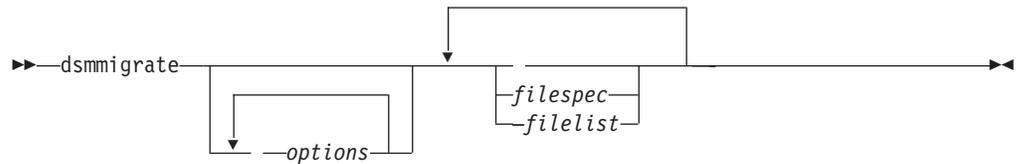
If the file resides in a file system with a different server stanza than the last processed file, a new session starts for each file that is migrated. This can happen as a result of links from one file system to another.

Note:

1. If you set the *tapeprompt* option to *yes* in your *dsm.opt* file, and the destination for your migrated file is a storage pool that consists of removable media such as tape, you are prompted either to wait for the medium to mount, or to skip the file.
2. The Tivoli Storage Manager for Space Management client does not migrate contents of symbolic links. Symbolic links are not followed during recursive selective migration (**dsmmigrate -r** command).

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 your file system. If the total number of bytes exceeds the quota after the file migrates, the next file is not migrated.

Syntax



Parameters

Options

Select any of these options:

-Premigrate or -p

Files that are migrated are in the premigrated state.

-Recursive or -r

Migrates files in any subdirectory below the specified directory that match the file specification. If you do not use this option, only those files from the directories that you specify are migrated.

See note 2 on page 124 for information about **dsmmigrate -r** and symbolic links.

-Detail or -v

Displays the size and file name for each file that you migrate.

filespec

The path and file name of the file 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. Or, 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

Processes a list of files.

Examples

Task	Command
Migrate all files in a directory named /migfs2/test/dir1 and in all of its subdirectories. Display the information.	<code>dsmmigrate -Recursive -Detail /migfs2/test/dir1</code>
Migrate a file named tf04 from the current directory and display the information.	<code>dsmmigrate -Detail tf04</code>
Migrate all files in a filelist named /tmp/filelist	<code>dsmmigrate -filelist=/tmp/filelist</code>

dsmmigundelete

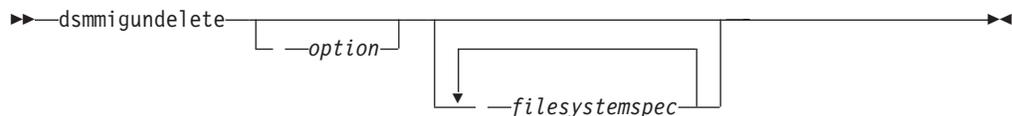
The **dsmmigundelete** command recreates deleted stub files for migrated files, and creates stub files for premigrated files for which an original file does not exist on your local file system. The file then becomes a migrated file.

You must have root user authority to use this command.

When you delete a stub file or an original copy of a premigrated file 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 hardlinked files. If you attempt to recreate a stub file for a hardlinked file, a stub file is not recreated *unless* all of the files that are hardlinked together are deleted from your local file system. When one file in a set of hardlinked files is migrated, all of the hardlinked files in the set become stub files. When the **dsmmigundelete** command recreates a stub file for a hardlinked 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 hardlinked set of files.

Syntax



Parameters

-Recover or **-r**

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 will overwrite 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 cannot check whether the data of the local stub file is different from the contents of the file in the TSM server.

Use this option only if all or most of the files in the file system are stub files and if you need good performance during the stub file recreation.

For performance purposes, do not use this option with the `-detail` option.

Files that were never migrated are not recreated.

-Detail or **-d**

Displays detailed information about which stub file is being recreated to the file system.

Do not use this option with the `-recover` option.

-Expiring or **-e**

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.

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 for migrated files that are marked for expiration that were accidentally deleted from the /home file system, and for files that are not marked for expiration. Reconciliation was run since the files were deleted.	<code>dsmmigundelete -expiring /home</code>
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.	<code>dsmmigundelete /home</code>
Recreates stub files for migrated files that had been removed from the /trul1ofs file system.	<code>dsmmigundelete -recover /trul1ofs</code>

dsmmonitord

The **dsmmonitord** 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`—▶▶

dsmq

The **dsmq** command displays following information about each recall process that is queued for processing.

You must have root user authority to use this command.

Specifically the **dsmq** command displays:

- The recall ID
- The hostname of the host recalling the file
- The start time for the recall process

- The inode number for the recalled file
- The name of the file system where the file is being recalled
- The original name of the file at the time it was migrated

If you set the *maxrecalldaemons* option in your *dsm.sys* file lower than the current number of requested recalls, some recall requests will not appear in the output for this command until recall daemons are available to perform the requests. To remove a recall process from the queue, use the **dsmrm** command.

Note: If the Tivoli Storage Manager server is busy, or the connection between the HSM client and the Tivoli Storage Manager server is slow, the original name of the file might display as UNKNOWN in the output for this command. Issue **dsmq** again to view the file name.

If the recall daemon process ID (DPID) is zero, the recall is complete. You cannot remove the recall process from the queue.

Syntax

▶▶ `dsmq` ◀◀

Examples

Task	Command
Display the status of recall processes.	<code>dsmq</code>

dsmrecall

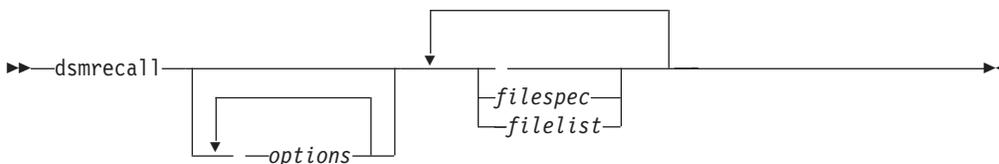
The **dsmrecall** command selectively recalls migrated files or parts of migrated files to your local file system. Space management must be active.

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 that of the last processed file, a new session is started for each file that is recalled. This can happen as a result of links from one file system to another.

To display a list of all your migrated files, use the **dsmmigquery** command. To display information about a list of migrated files from a particular file system or directory, use the **dsmils** command.

Syntax



Parameters

Options

-Recursive or -r

Recalls 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.

-Detail or -v

Displays the size, path, and file name for each file that you recall.

-OFFSET=n or -on

Specifies the offset from the beginning of the file of the required data range for partial recall (in bytes, kilobytes, megabytes, or gigabytes). You can use multipliers (k, m, g, K, M, or G) can be used. The range of values is 0 through 2147483647. There is no default value. This option is valid on AIX GPFS and Linux x86/x86_64 GPFS file systems only.

Remember: Use this option only with the *-size* option and the *filespec*. The *-recursive*, *-detail*, and *-filelist* options are not valid with the *-offset* option.

Examples are: *-offset=10* (bytes), *-offset=23k* (kilobytes), *-o5M* (megabytes), *-o2G* (gigabytes).

-SIZE=n or -sn

Specifies 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 values is 0 to 4294967295. There is no default value. This option is valid on AIX GPFS and Linux x86/x86_64 GPFS file systems only.

Remember: Use this option only with the *-offset* and the *filespec*. The *-recursive*, *-detail*, and *-filelist* options are not valid with the *-size* option.

Examples are: *-size=10* (bytes), *-size=23k* (kilobytes), *-s5M* (megabytes), *-s2G* (gigabytes).

filespec

The path and file name of the file that you want to recall. This parameter is required. You can use wildcard characters to specify a group of files or all files in a directory, or you can use more than one file specification in one command. When you use wildcard characters in a file specification, the HSM client attempts to recall 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

Processes a list of files.

Examples

Task	Command
Recall a single file named <i>/migfs1/test/tf04</i> and display detailed information.	<code>dsmrecall -Detail /migfs1/test/tf04</code>
Recall all migrated files in a directory named <i>/mfs4/user1</i> and all migrated files in its subdirectories.	<code>dsmrecall -Recursive /mfs4/user1/*</code>
Recall all files in the filelist named <i>/tmp/filelist</i>	<code>dsmrecall -filelist=/tmp/filelist</code>
Recall the specified portion of a file named <i>/mfs1/file10</i> in partial recall mode.	<code>dsmrecall -offset=10M -size=500M /mfs1/file10</code>

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 HSM managed files, **dsmrecalld** must be running.

Syntax

►—dsmrecalld—◄

dsmreconcile

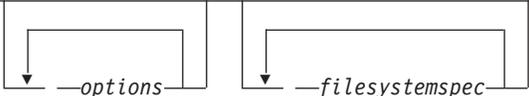
The **dsmreconcile** command synchronizes the file systems on your client node with the Tivoli Storage Manager server that you contact for space management services. Use this command at any time to perform some or all of the reconciliation tasks for one or more file systems. Specify how often to perform automatic reconciliation 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 perform reconciliation.

Syntax

►—dsmreconcile—◄



Parameters

Options

—detail or -d

Prints progress messages while receiving the list of migrated files from the server.

—fileinfo or -f

Performs the basic reconciliation task of expiring and deleting obsolete copies from the server. This is the default.

—orphancheck or -o

Checks for orphaned files in the local file system.

The orphancheck reconciliation process can take several minutes to process depending on the number of migrated objects and the size of the CFI database.

The HSM client verifies that valid objects for each local stub file exists in the space management pool of the Tivoli Storage Manager server. The

dsmreconcile process queries the scout daemon for all migrated and premigrated files and checks that corresponding objects exist on the server. When orphans are located, their names are recorded in the `.SpaceMan/orphan.stubs` file. This option overrides the `no` value that you set for the `checkfororphans` option in the `dsm.sys` file.

Requirement: Because the scout daemon also aids in the reconciliation process, it must be running.

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.

filesystemspec

The name of the file system to reconcile. If you do not specify a file system name, the HSM 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 -o /migfs1</code>
Start reconciliation for file systems <code>/home</code> and <code>/test1</code> .	<code>dsmreconcile /home /test1</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.

Examples

Task	Command
Remove recall ID 10 from the queue.	dsmrm 10
Remove recall ID 5 and recall ID 6 from the queue.	dsmrm 5 6

dsmrootd

The **dsmrootd** command starts a root daemon if it is not running. The root daemon provides non-root user support for the HSM client.

You must have root user authority to use this command.

The following commands can be run as non-root:

- **dsmdf**
- **dsmdu**
- **dsmls**
- **dsmmigfs query**
- **dsmmighelp**
- **dsmmigrate**
- **dsmrecall**

Note: If you issue this command while a root daemon is running, no action is taken.

Syntax

▶▶ dsmrootd ◀◀

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

▶▶ dsmscoutd ◀◀

start
stop
restart
scanplan
all
<filesystemspec>
scannow
<filesystemspec>
all
statistics
<filesystemspec>
help

Parameters

Start

Starts the scout daemon. This 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 HSM managed file systems, the `dsmscoutd` daemon lost its communication paths due to unexpected file system issues such as a file system being unmounted. 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 to perform the specified action. 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.

dsmsetpw

The `dsmsetpw` command changes the Tivoli Storage Manager 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 you did not register your client node with the Tivoli Storage Manager server that you contact for services, and open registration is in effect, the Tivoli Storage Manager client prompts you for registration information.

Syntax

```
▶▶ dsmsetpw [ -oldpassword -newpassword ] ▶▶
```

Parameters

oldpassword

The current Tivoli Storage Manager password for your client node.

newpassword

The new Tivoli Storage Manager password to set for your client node.

Examples

Task	Command
Change your current Tivoli Storage Manager password from <i>osecret</i> to <i>nsecret</i> .	<code>dsmsetpw osecret nsecret</code>

dsmwatchd

The **dsmwatchd** command manages failover activities for your HSM client node. If failover is active on your HSM client node, it checks the status of the **dsmrecalld** command, the **dsmmonitor** command, the **dsmscoutd** command, and the **dsmrootd** command. If any of these daemons end, or becomes corrupted, **dsmwatchd** automatically recovers the failed daemon.

Valid on AIX GPFS and Linux x86/x86_64 GPFS file systems only.

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.

If two or more nodes within a GPFS node set participate actively in a failover environment, the **dsmwatchd** command either will take over the file systems of a failed HSM node actively (remote is the same as within the local nodeset), or start the failover if the HSM client no longer can perform its operations locally. A node crash also can start failover. Unmounting a managed file system will not result in a failover. The failover environment is active by default. Issue the **dsmmigfs ENABLEFailover** or **dsmmigfs DISABLEFailover** commands to change the status.

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:

1. 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.
2. Set the environment variable, `DSM_LOG`. For example:
`DSM_LOG=/usr/tivoli.tsm/client/hsm/bin/dsmerror.log`

Note:

1. For AIX GPFS, set the `DSM_LOG` in the `/etc/environment` file accordingly.
2. For Linux x86/x86_64 GPFS, the `/etc/environment` file does not exist. Use the *errorlogname* option.

Syntax

▶—dsmwatchd—▶

Appendix A. The SpaceMan directory

Tivoli Storage Manager creates file system-specific control files that reside in a hidden directory named SpaceMan on each file system to which you add space management.

These files also reside in the `/etc/adsm/SpaceMan` directory that Tivoli Storage Manager creates when you install the HSM client. The HSM client automatically excludes these files from space management. User action is not required to ensure that the files remain on your local file systems. Table 30 describes the file contents of this directory.

Table 30. HSM Control Files

File Name	Description
<code>config/dmiFSGlobalState</code>	Valid for AIX GPFS, Linux x86/x86_64 GPFS, Solaris VxFS, and HP-UX VxFS file systems <i>only</i> . Global and HSM-managed file system information. Related to the DMAPI interface.
<code>config/dmiFSGlobalState.pid</code>	A lock file for the <code>dmiFSGlobalState</code> file.
<code>ActiveRecallTab</code>	Active recall table.
<code>config/DSM.pid</code>	A lock file for the <code>DSMNodeSet</code> and <code>DSMSDRVersion</code> files.
<code>config/DSMNodeSet</code>	A file storing the HSM configuration.
<code>config/DSMSDRVersion</code>	A file storing the number of changes in the HSM configuration.

For HSM storage pool support, all files located in SpaceMan are placed in the same storage pool. To achieve this and to prevent other migration rules from moving those files to a different pool, there are EXCLUDE rules that help achieve these goals.

To ensure that other migration rules do not move files into Spaceman, you can specify:

```
RULE 'TSM_EXCL_DOTSPACEMAN' EXCLUDE WHERE PATH_NAME LIKE '%/.SpaceMan/%'
```

If files were moved to an unwanted location, the following example contains a migration rule that you can use to move 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.

Appendix B. HSM .SpaceMan files and subdirectories

When you add space management to your file systems, the HSM client creates control files in a hidden directory named `.SpaceMan`.

Table 31 provides a brief description of objects that are stored in the `.SpaceMan` directory in each file system to which you add space management. The HSM client requires these objects for processing.

Attention: The HSM client can create other objects in the `.SpaceMan` directory during space management processing. *Do not* delete or modify the `.SpaceMan` directory, its contents, or any of the file ownerships and permissions.

Table 31. Control files stored in the `.SpaceMan` directory

Files	Description
<code>orphan.stubs</code>	Records the names of any files for which a stub file exists on your local file system, but a corresponding migrated file does not exist in Tivoli Storage Manager storage.
<code>status</code>	Records space management-related statistics for your file system.
<code>dmiFSState</code>	Stores information about the file system.
<code>dsmmigfstab</code>	If this file exists, it is from an older version of the HSM client. It will not be modified nor deleted, but is needed for transition to current version.
<code>hsmfsconfig.pid</code>	A lock file for the local <code>hsmfsconfig.xml</code> file.
<code>hsmfsconfig.xml</code>	A file containing the space management settings that you selected for this file system.
<code>logdir/</code>	Records information during file migration or recall. It uses this information to complete any interrupted transactions, such as a system failure.
<code>SDR/</code>	Valid for AIX GPFS and Linux x86/x86_64 GPFS only If this directory exists, it is from an older version of the HSM client.
<code>reserved/</code>	Contains reserved files for demand migration.
<code>metadata/</code>	This directory contains the CFI that was created by the scout daemon. The CFI contains file system information.

Appendix C. Accessibility features for Tivoli Storage Manager

Accessibility features help a user who has a physical disability, such as restricted mobility or limited vision, to use software products successfully. The major accessibility features of Tivoli Storage Manager are described in this topic.

Accessibility features

The following list includes the major accessibility features in Tivoli Storage Manager:

- Keyboard-only operation
- Interfaces that are commonly used by screen readers
- Keys that are discernible by touch but do not activate just by touching them
- Industry-standard devices for ports and connectors
- The attachment of alternative input and output devices
- User documentation provided in HTML and PDF format. Descriptive text is provided for all documentation images.

The Tivoli Storage Manager Information Center, and its related publications, are accessibility-enabled.

Keyboard navigation

The Tivoli Storage Manager for Windows Console follows Microsoft conventions for all keyboard navigation and access. Drag and Drop support is managed using the Microsoft Windows Accessibility option known as MouseKeys. For more information about MouseKeys and other Windows accessibility options, please refer to the Windows Online Help (keyword: MouseKeys).

Tivoli Storage Manager follows AIX operating system conventions for keyboard navigation and access.

Tivoli Storage Manager follows HP-UX operating-system conventions for keyboard navigation and access.

Tivoli Storage Manager follows Linux operating-system conventions for keyboard navigation and access.

Tivoli Storage Manager follows Sun Solaris operating-system conventions for keyboard navigation and access.

Tivoli Storage Manager follows z/OS[®] operating-system conventions for keyboard navigation and access.

Vendor software

Tivoli Storage Manager 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 the accessibility information about its products.

Related accessibility information

You can view the publications for Tivoli Storage Manager in Adobe® Portable Document Format (PDF) using the Adobe Acrobat Reader. You can access these or any of the other documentation PDFs at the IBM Publications Center at <http://www.ibm.com/shop/publications/order/>.

IBM and accessibility

For more information about the commitment that IBM has to accessibility, see the IBM Human Ability and Accessibility Center at <http://www.ibm.com/able>.

Notices

This information was developed for products and services offered in the U.S.A.

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
Armonk, NY 10504-1785
U.S.A.*

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:

*IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106-0032, Japan*

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

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 states 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 Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites 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 Corporation
2Z4A/101
11400 Burnet Road
Austin, TX 78758
U.S.A.*

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 information 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.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com) are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at <http://www.ibm.com/legal/copytrade.shtml>.

Java and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

Glossary

This glossary includes terms and definitions for IBM Tivoli Storage Manager.

To view glossaries for other IBM products, go to <http://www.ibm.com/software/globalization/terminology/>.

The following cross-references are used in this glossary:

- *See* refers the reader from a term to a preferred synonym, or from an acronym or abbreviation to the defined full form.
- *See also* refers the reader to a related or contrasting term.

A

absolute mode

In storage management, a backup copy-group mode that specifies that a file is considered for incremental backup even if the file has not changed since the last backup. See also *modified mode*.

access control list (ACL)

In computer security, a list associated with an object that identifies all the subjects that can access the object and their access rights. For example, an access control list is associated with a file that identifies the users who can access that file and their access rights.

access mode

An attribute of a storage pool or a storage volume that specifies whether the server can write to or read from the storage pool or storage volume. The access mode can be read/write, read-only, or unavailable. Volumes in primary storage pools can also have an access mode of destroyed. Volumes in copy storage pools can also have an access mode of offsite.

acknowledgment

The transmission of acknowledgment characters as a positive response to a data transmission.

ACL See *access control list*.

activate

To validate the contents of a policy set and then make it the active policy set.

active-data pool

A named set of storage pool volumes that contain only active versions of client backup data.

active file system

A file system to which space management has been added. With space management, tasks for an active file system include automatic migration, reconciliation, selective migration, and recall. Contrast with *inactive file system*.

active policy set

The activated policy set that contains the policy rules in use by all client nodes that are assigned to the policy domain. See also *policy domain* and *policy set*.

active version

The most recent backup copy of a file stored. The active version of a file cannot be deleted until a backup process detects that the user has either replaced the file with a newer version or has deleted the file from the file server or workstation. Contrast with *inactive version*.

activity log

A log that records normal activity messages that are generated by the server. These messages include information about server and client operations, such as the start time of sessions or device I/O errors.

adaptive subfile backup

A type of backup that sends only changed portions of a file to the server, instead of sending the entire file. Adaptive subfile backup reduces network traffic and increases the speed of the backup.

administrative client

A program that runs on a file server, workstation, or mainframe that administrators use to control and monitor the Tivoli Storage Manager server. Contrast with *backup-archive client*.

administrative command schedule

A database record that describes the planned processing of an administrative

command during a specific time period.
See also *client schedule*.

administrative privilege class

See *privilege class*.

administrative session

A period of time during which an administrator user ID communicates with a server to perform administrative tasks. Contrast with *client node session*.

administrator

A user who is registered to the server as an administrator, and who is authorized to perform tasks and issue commands through the assignment of an administrative privilege class.

Advanced Program-to-Program Communication (APPC)

An implementation of the SNA LU 6.2 protocol that allows interconnected systems to communicate and share the processing of programs.

agent node

A client node that has been granted proxy authority to perform operations on behalf of another client node, which is the target node.

aggregate

An object, stored in one or more storage pools, consisting of a group of logical files that are packaged together. See also *logical file* and *physical file*.

aggregate data transfer rate

A performance statistic that indicates the average number of bytes that were transferred per second while processing a given operation.

APPC See *Advanced Program-to-Program Communication*.

application client

A program that is installed on a system to protect an application. The Tivoli Storage Manager server provides backup services to an application client.

archive

To copy programs, data, or files to another storage media, usually for long-term storage or security. Contrast with *retrieve*.

archive copy

A file or group of files that was archived to server storage.

archive copy group

A policy object containing attributes that control the generation, destination, and expiration of archived files.

archive-retention grace period

The number of days that the storage manager retains an archived file when the server is unable to rebind the file to an appropriate management class. See also *bind*.

association

(1) The defined relationship between a client node and a client schedule. An association identifies the name of a schedule, the name of the policy domain to which the schedule belongs, and the name of a client node that performs scheduled operations.

(2) On a configuration manager, the defined relationship between a profile and an object such as a policy domain. Profile associations define the configuration information that is distributed to a managed server when it subscribes to the profile.

audit

To check for logical inconsistencies between information that the server has and the actual condition of the system. The storage manager can audit information about items such as volumes, libraries, and licenses. For example, when a storage manager audits a volume, the server checks for inconsistencies between information about backed-up or archived files that are stored in the database and the actual data that are associated with each backup version or archive copy in server storage.

authentication

The process of checking a user's password before permitting user access to the Tivoli Storage Manager server. Authentication can be turned on or off by an administrator with system privilege.

authentication rule

A specification that another user can use to either restore or retrieve files from storage.

authority

The right to access objects, resources, or functions. See also *privilege class*.

authorization rule

A specification that permits another user to either restore or retrieve a user's files from storage.

authorized user

A user who has administrative authority for the Tivoli Storage Manager client on a workstation. This user changes passwords, performs open registrations, and deletes file spaces.

AutoFS

See *automounted file system*.

automatic detection

A feature that detects, reports, and updates the serial number of a drive or library in the database when the path from the local server is defined.

automatic migration

The process that is used to automatically move files from a local file system to storage, based on options and settings that are chosen by a root user on a workstation. See also *threshold migration* and *demand migration*.

automatic reconciliation

The process that is used to reconcile file systems at regular intervals. The intervals are set by a user with root user authority. See also *reconciliation*.

automounted file system (AutoFS)

A file system that is managed by an automounter daemon. The automounter daemon monitors a specified directory path, and automatically mounts the file system to access data.

B**backup-archive client**

A program that runs on a workstation or file server and provides a means for users to back up, archive, restore, and retrieve files. Contrast with *administrative client*.

backup copy group

A policy object containing attributes that control the generation, destination, and expiration of backup versions of files. A backup copy group belongs to a management class.

backup-retention grace period

The number of days the storage manager retains a backup version after the server is unable to rebind the file to an appropriate management class.

backup set

A portable, consolidated group of active versions of backup files that are generated for a backup-archive client.

backup set collection

A group of backup sets that are created at the same time and which have the same backup set name, volume names, description, and device classes. The server identifies each backup set in the collection by its node name, backup set name, and file type.

backup version

A file or directory that a client node backed up to server storage. More than one backup version can exist in server storage, but only one backup version is the active version. See also *active version* and *inactive version*.

bindery

A database that consists of three system files for a NetWare server. The files contain user IDs and user restrictions.

bind

To associate a file with a management class name. See *rebind*.

C**cache**

To place a duplicate copy of a file on random access media when the server migrates a file to another storage pool in the hierarchy.

cache file

A snapshot of a logical volume created by Logical Volume Snapshot Agent. Blocks are saved immediately before they are modified during the image backup and their logical extents are saved in the cache files.

CAD

See *client acceptor*.

central scheduler

A function that permits an administrator to schedule client operations and administrative commands. The operations can be scheduled to occur periodically or on a specific date. See *client schedule* and *administrative command schedule*.

client A software program or computer that requests services from a server.

client acceptor daemon (CAD)

See *client acceptor*.

client acceptor

An HTTP service that serves the Java applet for the Web client to Web browsers. On Windows systems, the client acceptor is installed and run as a service. On AIX, UNIX, and Linux systems, the client acceptor is run as a daemon, and is also called the *client acceptor daemon* (CAD).

client domain

The set of drives, file systems, or volumes that the user selects to back up or archive data, using the backup-archive client.

client node

A file server or workstation on which the backup-archive client program has been installed, and which has been registered to the server.

client node session

A session in which a client node communicates with a server to perform backup, restore, archive, retrieve, migrate, or recall requests. Contrast with *administrative session*.

client options file

An editable file that identifies the server and communication method, and provides the configuration for backup, archive, hierarchical storage management, and scheduling.

client option set

A group of options that are defined on the server and used on client nodes in conjunction with client options files.

client-polling scheduling mode

A method of operation in which the client queries the server for work. Contrast with *server-prompted scheduling mode*.

client schedule

A database record that describes the planned processing of a client operation during a specific time period. The client operation can be a backup, archive, restore, or retrieve operation, a client operating system command, or a macro. See also *administrative command schedule*.

client/server

Pertaining to the model of interaction in distributed data processing in which a program on one computer sends a request to a program on another computer and awaits a response. The requesting program is called a client; the answering program is called a server.

client system-options file

A file, used on AIX, UNIX, or Linux system clients, containing a set of processing options that identify the servers to be contacted for services. This file also specifies communication methods and options for backup, archive, hierarchical storage management, and scheduling. This file is also called the *dsm.sys* file. See also *client user-options file*.

client user-options file

A file that contains the set of processing options that the clients on the system use. The set can include options that determine the server that the client contacts, and options that affect backup operations, archive operations, hierarchical storage management operations, and scheduled operations. This file is also called the *dsm.opt* file. For AIX, UNIX, or Linux systems, see also *client system-options file*.

closed registration

A registration process in which only an administrator can register workstations as client nodes with the server. Contrast with *open registration*.

collocation

The process of keeping all data belonging to a single-client file space, a single client node, or a group of client nodes on a minimal number of sequential-access volumes within a storage pool. Collocation can reduce the number of volumes that must be accessed when a large amount of data must be restored.

collocation group

A user-defined group of client nodes whose data is stored on a minimal number of volumes through the process of collocation.

commit point

A point in time when data is considered consistent.

communication method

The method by which a client and server exchange information. See also *Transmission Control Protocol/Internet Protocol*.

Common Programming Interface for Communications (CPI-C)

A call-level interface that provides a consistent application programming interface (API) for applications that use program-to-program communications. CPI-C uses LU 6.2 architecture to create a set of interprogram services that can establish and end a conversation, send and receive data, exchange control information, and notify a partner program of errors.

communication protocol

A set of defined interfaces that permit computers to communicate with each other.

compression

A function that removes repetitive characters, spaces, or strings of characters from the data being processed and replaces the repetitive characters with control characters. Compression reduces the amount of storage space that is required for the data.

configuration manager

A server that distributes configuration information, such as policies and schedules, to managed servers according to their profiles. Configuration information can include policy and schedules. See also *managed server* and *profile*.

conversation

A connection between two programs over a session that allows them to communicate with each other while processing a transaction.

copy backup

A full backup in which the transaction log files are not deleted so that backup procedures that use incremental or differential backups are not disrupted

copy group

A policy object containing attributes that control how backup versions or archive copies are generated, where backup versions or archive copies are initially

located, and when backup versions or archive copies expire. A copy group belongs to a management class. See also *archive copy group*, *backup copy group*, *backup version*, and *management class*.

copy storage pool

A named set of volumes that contain copies of files that reside in primary storage pools. Copy storage pools are used only to back up the data that is stored in primary storage pools. A copy storage pool cannot be a destination for a backup copy group, an archive copy group, or a management class (for space-managed files). See also *primary storage pool* and *destination*.

CPI-C See *Common Programming Interface for Communications*.

D**daemon**

A program that runs unattended to perform continuous or periodic functions, such as network control.

damaged file

A physical file in which Tivoli Storage Manager has detected read errors.

data access control mode

A mode that controls whether a command can access a migrated file, see a migrated file as zero-length, or receive an input/output error if it attempts to access a migrated file. See also *execution mode*.

database backup series

One full backup of the database, plus up to 32 incremental backups made since that full backup. Each full backup that is run starts a new database backup series. A number identifies each backup series.

database snapshot

A complete backup of the entire database to media that can be taken off-site. When a database snapshot is created, the current database backup series is not interrupted. A database snapshot cannot have incremental database backups associated with it. See also *database backup series*. Contrast with *full backup*.

data deduplication

A method of reducing storage needs by eliminating redundant data. Only one instance of the data is retained on storage

media, such as disk or tape. Other instances of the same data are replaced with a pointer to the retained instance.

data manager server

A server that collects metadata information for client inventory and manages transactions for the storage agent over the local area network. The data manager server informs the storage agent with applicable library attributes and the target volume identifier.

data mover

A device that moves data on behalf of the server. A network-attached storage (NAS) file server is a data mover.

data storage-management application-program interface (DSMAPI)

A set of functions and semantics that can monitor events on files, and manage and maintain the data in a file. In an HSM environment, a DSMAPI uses events to notify data management applications about operations on files, stores arbitrary attribute information with a file, supports managed regions in a file, and uses DSMAPI access rights to control access to a file object.

default management class

A management class that is assigned to a policy set. This class is used to govern backed up or archived files when a file is not explicitly associated with a specific management class through the include-exclude list.

deduplication

See *data deduplication*.

demand migration

The process that is used to respond to an out-of-space condition on a file system for which hierarchical storage management (HSM) is active. Files are migrated to server storage until space usage drops to the low threshold that was set for the file system. If the high threshold and low threshold are the same, one file is migrated.

desktop client

The group of backup-archive clients that includes clients on Microsoft Windows, Apple, and Novell NetWare operating systems.

destination

A copy group or management class attribute that specifies the primary storage pool to which a client file will be backed up, archived, or migrated.

device class

A named set of characteristics that are applied to a group of storage devices. Each device class has a unique name and represents a device type of disk, file, optical disk, or tape.

device configuration file

(1) For a server, a file that contains information about defined device classes, and, on some servers, defined libraries and drives. The information is a copy of the device configuration information in the database.

(2) For a storage agent, a file that contains the name and password of the storage agent, and information about the server that is managing the SAN-attached libraries and drives that the storage agent uses.

device driver

A program that provides an interface between a specific device and the application program that uses the device.

disaster recovery manager (DRM)

A function that assists in preparing and using a disaster recovery plan file for the server.

disaster recovery plan

A file that is created by the disaster recovery manager (DRM) that contains information about how to recover computer systems if a disaster occurs and scripts that can be run to perform some recovery tasks. The file includes information about the software and hardware that is used by the server, and the location of recovery media.

domain

A grouping of client nodes with one or more policy sets, which manage data or storage resources for the client nodes. See *policy domain* or *client domain*.

DRM See *disaster recovery manager*.

DSMAPI

See *data storage-management application-program interface*.

dynamic serialization

A type of copy serialization in which a file or folder is backed up or archived on the first attempt regardless of whether it changes during a backup or archive.

E

EA See *extended attribute*.

EB See *exabyte*.

EFS See *Encrypted File System*.

Encrypted File System (EFS)

A file system that uses file system-level encryption.

enterprise configuration

A method of setting up servers so that the administrator can distribute the configuration of one of the servers to the other servers, using server-to-server communication. See also *configuration manager*, *managed server*, *profile*, and *subscription*.

enterprise logging

The process of sending events from a Tivoli Storage Manager server to a designated event server. The event server routes the events to designated receivers, such as to a user exit. See also *event*.

error log

A data set or file that is used to record error information about a product or system.

estimated capacity

The available space, in megabytes, of a storage pool.

event (1) An administrative command or a client operation that is scheduled to be run using Tivoli Storage Manager scheduling.

(2) A message that an Tivoli Storage Manager server or client issues. Messages can be logged using Tivoli Storage Manager event logging.

event record

A database record that describes actual status and results for events.

event server

A server to which other servers can send events for logging. The event server routes the events to any receivers that are enabled for the sending server's events.

exabyte (EB)

For processor storage, real and virtual storage, and channel volume, 1 152 921 504 606 846 976 bytes. For disk storage capacity and communications volume, 1 000 000 000 000 000 000 bytes.

exclude

The process of identifying files in an include-exclude list. This process prevents the files from being backed up or migrated whenever a user or schedule enters an incremental or selective backup operation. A file can be excluded from backup and space management, backup only, or space management only.

exclude-include list

See *include-exclude list*.

execution mode

A mode that controls the space-management related behavior of commands that run under the **dsmmode** command.

expiration

The process by which files, data sets, or objects are identified for deletion because their expiration date or retention period has passed.

expiring file

A migrated or premigrated file that has been marked for expiration and removal from storage. If a stub file or an original copy of a premigrated file is deleted from a local file system, or if the original copy of a premigrated file is updated, the corresponding migrated or premigrated file is marked for expiration the next time reconciliation is run.

extend

To increase the portion of available space that can be used to store database or recovery log information.

extended attribute (EA)

Names or value pairs that are associated with files or directories. There are three classes of extended attributes: user attributes, system attributes, and trusted attributes.

external library

A type of library that is provided by Tivoli Storage Manager that permits LAN-free data movement for StorageTek libraries that are managed by Automated

Cartridge System Library Software (ACSLs). To activate this function, the Tivoli Storage Manager library type must be EXTERNAL.

F

file access time

On AIX, UNIX, or Linux systems, the time when the file was last accessed.

file age

For migration prioritization purposes, the number of days since a file was last accessed.

file device type

A device type that specifies the use of sequential access files on disk storage as volumes.

file server

A dedicated computer and its peripheral storage devices that are connected to a local area network that stores programs and files that are shared by users on the network.

file space

A logical space in server storage that contains a group of files that have been backed up or archived by a client node, from a single logical partition, file system, or virtual mount point. Client nodes can restore, retrieve, or delete their file spaces from server storage. In server storage, files belonging to a single file space are not necessarily stored together.

file space ID (FSID)

A unique numeric identifier that the server assigns to a file space when it is stored in server storage.

file state

The space management mode of a file that resides in a file system to which space management has been added. A file can be in one of three states: resident, premigrated, or migrated. See also *resident file*, *premigrated file*, and *migrated file*.

file system migrator (FSM)

A kernel extension that intercepts all file system operations and provides any space management support that is required. If no space management support is required, the operation is passed to the operating system, which performs its normal functions. The file system

migrator is mounted over a file system when space management is added to the file system.

file system state

The storage management mode of a file system that resides on a workstation on which the hierarchical storage management (HSM) client is installed. A file system can be in one of these states: native, active, inactive, or global inactive.

frequency

A copy group attribute that specifies the minimum interval, in days, between incremental backups.

FSID See *file space ID*.

FSM See *file system migrator*.

full backup

The process of backing up the entire server database. A full backup begins a new database backup series. See also *database backup series* and *incremental backup*. Contrast with *database snapshot*.

fuzzy backup

A backup version of a file that might not accurately reflect what is currently in the file because the file was backed up at the same time as it was being modified.

fuzzy copy

A backup version or archive copy of a file that might not accurately reflect the original contents of the file because it was backed up or archived the file while the file was being modified. See also *backup version* and *archive copy*.

G

General Parallel File System

A high-performance shared-disk file system that can provide data access from nodes in a cluster environment.

gigabyte (GB)

In decimal notation, 1 073 741 824 when referring to memory capacity; in all other cases, it is defined as 1 000 000 000.

global inactive state

The state of all file systems to which space management has been added when space management is globally deactivated for a client node. When space management is globally deactivated, hierarchical storage management (HSM)

cannot perform migration, recall, or reconciliation. However, a root user can update space management settings and add space management to additional file systems. Users can access resident and premigrated files.

Globally Unique Identifier (GUID)

An algorithmically determined number that uniquely identifies an entity within a system.

GPFS See *General Parallel File System*.

GPFS node set

A mounted, defined group of GPFS file systems.

group backup

The backup of a group containing a list of files from one or more file space origins.

GUID See *Globally Unique Identifier*.

H

hierarchical storage management (HSM)

A function that automatically distributes and manages data on disk, tape, or both by regarding devices of these types and potentially others as levels in a storage hierarchy that range from fast, expensive devices to slower, cheaper, and possibly removable devices. The objectives are to minimize access time to data and maximize available media capacity.

hierarchical storage management (HSM) client

A client program that works with the Tivoli Storage Manager server to provide hierarchical storage management (HSM) for a system. See also *hierarchical storage management* and *space manager client*.

HSM See *hierarchical storage management*.

HSM client

See *hierarchical storage management client*.

I

ILM See *information lifecycle management*.

image A file system or raw logical volume that is backed up as a single object.

image backup

A backup of a full file system or raw logical volume as a single object.

inactive file system

A file system for which space

management has been deactivated. Contrast with *active file system*.

inactive version

A backup version of a file that is either not the most recent backup version, or that is a backup version of a file that no longer exists on the client system. Inactive backup versions are eligible for expiration processing according to the management class assigned to the file. Contrast with *active version*.

include-exclude file

A file containing statements to determine the files to back up and the associated management classes to use for backup or archive. See also *include-exclude list*.

include-exclude list

A list of options that include or exclude selected files for backup. An exclude option identifies files that should not be backed up. An include option identifies files that are exempt from the exclusion rules or assigns a management class to a file or a group of files for backup or archive services.

incremental backup

(1) A copy of all database data that has changed since the most recent successful full backup operation. An incremental backup is also known as a *cumulative backup image* because each incremental backup includes the contents of the previous incremental backup.

(2) The process of backing up information in the database that is new or changed since the last full backup. Contrast with *full backup*. See also *database backup series*.

(3) For Data Protection for Microsoft Exchange Server, a backup in which the transaction logs are backed up and then cleared.

individual mailbox restore

See *mailbox restore*.

information lifecycle management (ILM)

GPFS policy-based file management for storage pools and file sets.

i-node The internal structure that describes the individual files on AIX, UNIX, or Linux systems. An i-node contains the node, type, owner, and location of a file.

i-node number

A number specifying a particular i-node file in the file system.

IP address

A unique address for a device or logical unit on a network that uses the IP standard.

J**job file**

A generated file that contains configuration information for a migration job. The file is XML format and can be created and edited in the hierarchical storage management (HSM) client for Windows client graphical user interface.

journal-based backup

A method for backing up Windows clients and AIX clients that exploits the change notification mechanism in a file to improve incremental backup performance by reducing the need to fully scan the file system.

journal daemon

On AIX, UNIX, or Linux systems, a program that tracks change activity for files residing in file systems.

journal service

In Microsoft Windows, a program that tracks change activity for files residing in file systems.

K**kilobyte (KB)**

For processor storage, real and virtual storage, and channel volume, 210 or 1 024 bytes. For disk storage capacity and communications volume, 1 000 bytes.

L

LAN See *local area network*.

LAN-free data movement

The movement of client data between a client system and a storage device on a storage area network (SAN), bypassing the local area network. This process is also referred to as *LAN-free data transfer*.

LAN-free data transfer

See *LAN-free data movement*.

leader data

Bytes of data, from the beginning of a migrated file, that are stored in the file's

corresponding stub file on the local file system. The amount of leader data that is stored in a stub file depends on the stub size that is specified.

library

(1) A repository for demountable recorded media, such as magnetic disks and magnetic tapes.

(2) A collection of one or more drives, and possibly robotic devices (depending on the library type), which can be used to access storage volumes.

library client

A server that uses server-to-server communication to access a library that is managed by another storage management server. See also *library manager*.

library manager

A server that controls device operations when multiple storage management servers share a storage device. See also *library client*.

local Pertaining to a device, file, or system that is accessed directly from a user's system, without the use of a communication line.

local area network (LAN)

A network that connects several devices in a limited area (such as a single building or campus) and that can be connected to a larger network.

local shadow volumes

Data that is stored on shadow volumes localized to a disk storage subsystem.

LOFS See *loopback virtual file system*.

logical file

A file that is stored in one or more server storage pools, either by itself or as part of an aggregate. See also *aggregate* and *physical file*.

logical occupancy

The space that is used by logical files in a storage pool. This space does not include the unused space created when logical files are deleted from aggregate files, so it might be less than the physical occupancy.

logical unit (LU)

An access point through which a user or application program accesses the Systems

Network Architecture (SNA) network to communicate with another user or application program.

logical unit number (LUN)

In the Small Computer System Interface (SCSI) standard, a unique identifier that is used to differentiate devices, each of which is a logical unit (LU).

logical volume

A portion of a physical volume that contains a file system.

logical volume backup

A back up of a file system or logical volume as a single object.

Logical Volume Snapshot Agent (LVSA)

Software that can act as the snapshot provider for creating a snapshot of a logical volume during an online image backup.

loopback virtual file system (LOFS)

A file system that is created by mounting a directory over another local directory, also known as mount-over-mount. A LOFS can also be generated using an automounter.

LU See *logical unit*.

LUN See *logical unit number*.

LVSA See *Logical Volume Snapshot Agent*.

M

MB See *megabyte*.

macro file

A file that contains one or more storage manager administrative commands, which can be run only from an administrative client using the MACRO command. Contrast with *Tivoli Storage Manager command script*.

mailbox restore

A function that restores Microsoft Exchange Server data (from IBM Data Protection for Exchange backups) at the mailbox level or mailbox-item level.

managed object

In Tivoli Storage Manager, a definition in the database of a managed server that was distributed to the managed server by a configuration manager. When a managed server subscribes to a profile, all objects that are associated with that

profile become managed objects in the database of the managed server. In general, a managed object cannot be modified locally on the managed server. Objects can include policy, schedules, client option sets, server scripts, administrator registrations, and server and server group definitions.

managed server

A Tivoli Storage Manager server that receives configuration information from a configuration manager using a subscription to one or more profiles. Configuration information can include definitions of objects such as policy and schedules. See also *configuration manager*, *subscription*, and *profile*.

management class

A policy object that users can bind to each file to specify how the server manages the file. The management class can contain a backup copy group, an archive copy group, and space management attributes. See also *copy group*, *space manager client*, *bind*, and *rebind*.

maximum transmission unit

The largest possible unit of data that can be sent on a given physical medium in a single frame. For example, the maximum transmission unit for Ethernet is 1500 bytes.

megabyte (MB)

- (1) 1 048 576 bytes (two to the twentieth power) when used in this publication.
- (2) For processor storage, real and virtual storage, and channel volume, 2 to the power of 20 or 1 048 576 bits. For disk storage capacity and communications volume, 1 000 000 bits.

metadata

Data that describes the characteristics of data; descriptive data.

migrate

To move data from one storage location to another. In Tivoli Storage Manager products, migrating can mean moving data from a client node to server storage, or moving data from one storage pool to the next storage pool defined in the server storage hierarchy. In both cases the

movement is controlled by policy, such as thresholds that are set. See also *migration threshold*.

migrated file

A file that has been copied from a local file system to Tivoli Storage Manager storage. For HSM clients on UNIX or Linux systems, the file is replaced with a stub file on the local file system. On Windows systems, creation of the stub file is optional. See also *stub file* and *resident file*. For HSM clients on UNIX or Linux systems, contrast with *premigrated file*.

migrate-on-close recall mode

A mode that causes a migrated file to be recalled back to its originating file system temporarily. Contrast with *normal recall mode* and *read-without-recall recall mode*.

migration job

A specification of files to migrate, and actions to perform on the original files after migration. See also *job file*.

migration threshold

High and low capacities for storage pools or file systems, expressed as percentages, at which migration is set to start and stop.

mirroring

The process of writing the same data to multiple locations at the same time. Mirroring data protects against data loss within the recovery log.

mode A copy group attribute that specifies whether to back up a file that has not been modified since the last time the file was backed up. See *modified mode* and *absolute mode*.

modified mode

In storage management, a backup copy-group mode that specifies that a file is considered for incremental backup only if it has changed since the last backup. A file is considered a changed file if the date, size, owner, or permissions of the file have changed. See also *absolute mode*.

mount limit

The maximum number of volumes that can be simultaneously accessed from the same device class. The mount limit determines the maximum number of mount points. See also *mount point*.

mount point

On the Tivoli Storage Manager server, a logical drive through which volumes in a sequential access device class are accessed. For removable-media device types, such as tape, a mount point is a logical drive that is associated with a physical drive. For the file device type, a mount point is a logical drive that is associated with an I/O stream. The number of mount points for a device class is defined by the value of the mount limit attribute for that device class. See also *mount limit*.

mount retention period

The maximum number of minutes that the server retains a mounted sequential-access media volume that is not being used before it dismounts the sequential-access media volume.

mount wait period

The maximum number of minutes that the server waits for a sequential-access volume mount request to be satisfied before canceling the request.

MTU See *maximum transmission unit*.

N**Nagle algorithm**

An algorithm that reduces congestion of TCP/IP networks by combining smaller packets and sending them together.

named pipe

A type of interprocess communication that permits message data streams to pass between peer processes, such as between a client and a server.

NAS See *network-attached storage*.

NAS node

A client node that is a network-attached storage (NAS) file server. Data for the NAS node is transferred by a NAS file server that is controlled by the network data management protocol (NDMP). A NAS node is also called a NAS file server node.

native file system

A file system that is locally added to the file server and is not added for space management. The hierarchical storage

manager (HSM) client does not provide space management services to the file system.

native format

A format of data that is written to a storage pool directly by the Tivoli Storage Manager server. Contrast with *non-native data format*.

NDMP

See *Network Data Management Protocol*.

NetBIOS

See *Network Basic Input/Output System*.

network-attached storage (NAS) file server

A dedicated storage device with an operating system that is optimized for file-serving functions. A NAS file server can have the characteristics of both a node and a data mover.

Network Basic Input/Output System (NetBIOS)

A standard interface to networks and personal computers that is used on local area networks to provide message, print-server, and file-server functions. Application programs that use NetBIOS do not have to handle the details of LAN data link control (DLC) protocols.

Network Data Management Protocol (NDMP)

A protocol that allows a network storage-management application to control the backup and recovery of an NDMP-compliant file server, without installing vendor-acquired software on that file server.

network data-transfer rate

A rate that is calculated by dividing the total number of bytes that are transferred by the data transfer time. For example, this rate can be the time that is spent transferring data over a network.

node A file server or workstation on which the backup-archive client program has been installed, and which has been registered to the server.

node name

A unique name that is used to identify a workstation, file server, or PC to the server.

node privilege class

A privilege class that gives an administrator the authority to remotely access backup-archive clients for a specific

client node or for all clients in a policy domain. See also *privilege class*.

non-native data format

A format of data that is written to a storage pool that differs from the format that the server uses for operations.

normal recall mode

A mode that causes a migrated file to be copied back to its originating file system when it is accessed.

O

offline volume backup

A backup in which the volume is locked so that no other system applications can access it during the backup operation.

online volume backup

A backup in which the volume is available to other system applications during the backup operation.

open registration

A registration process in which users can register their workstations as client nodes with the server. Contrast with *closed registration*.

operator privilege class

A privilege class that gives an administrator the authority to disable or halt the server, enable the server, cancel server processes, and manage removable media. See also *privilege class*.

options file

A file that contains processing options. On Windows and NetWare systems, the file is called *dsm.opt*. On AIX, UNIX, Linux, and Mac OS X systems, the file is called *dsm.sys*.

originating file system

The file system from which a file was migrated. When a file is recalled using normal or migrate-on-close recall mode, it is always returned to its originating file system.

orphaned stub file

A file for which no migrated file can be found on the Tivoli Storage Manager server that the client node is contacting for space management services. For example, a stub file can be orphaned when the client system-options file is

modified to contact a server that is different than the one to which the file was migrated.

out-of-space protection mode

A mode that controls whether the program intercepts out-of-space conditions. See also *execution mode*.

P

pacing

In SNA, a technique by which the receiving system controls the rate of transmission of the sending system to prevent overrun.

packet In data communication, a sequence of binary digits, including data and control signals, that is transmitted and switched as a composite whole.

page A defined unit of space on a storage medium or within a database volume.

partial-file recall mode

A recall mode that causes the hierarchical storage management (HSM) function to read just a portion of a migrated file from storage, as requested by the application accessing the file.

password generation

A process that creates and stores a new password in an encrypted password file when the old password expires. Automatic generation of a password prevents password prompting. Password generation can be set in the options file (passwordaccess option). See also *options file*.

path An object that defines a one-to-one relationship between a source and a destination. Using the path, the source accesses the destination. Data can flow from the source to the destination, and back. An example of a source is a data mover (such as a network-attached storage [NAS] file server), and an example of a destination is a tape drive.

pattern-matching character

See *wildcard character*.

physical file

A file that is stored in one or more storage pools, consisting of either a single logical file, or a group of logical files that are packaged together as an aggregate. See also *aggregate* and *logical file*.

physical occupancy

The amount of space that is used by physical files in a storage pool. This space includes the unused space that is created when logical files are deleted from aggregates. See also *physical file*, *logical file*, and *logical occupancy*.

plug-in

A self-contained software component that modifies (adds, or changes) the function in a particular system. When a plug-in is added to a system, the foundation of the original system remains intact.

policy domain

A grouping of policy users with one or more policy sets, which manage data or storage resources for the users. The users are client nodes that are associated with the policy domain.

policy privilege class

A privilege class that gives an administrator the authority to manage policy objects, register client nodes, and schedule client operations for client nodes. Authority can be restricted to certain policy domains. See also *privilege class*.

policy set

A group of rules in a policy domain. The rules specify how data or storage resources are automatically managed for client nodes in the policy domain. Rules can be contained in management classes. See also *active policy set* and *management class*.

premigrated file

A file that has been copied to Tivoli Storage Manager storage, but has not been replaced with a stub file on the local file system. An identical copy of the file resides both on the local file system and in Tivoli Storage Manager storage. Premigrated files occur on UNIX and Linux file systems to which space management has been added. Contrast with *migrated file* and *resident file*.

premigrated files database

A database that contains information about each file that has been premigrated to Tivoli Storage Manager storage. The database is stored in a hidden directory

named `.SpaceMan` in each file system to which space management has been added.

premigration

The process of copying files that are eligible for migration to Tivoli Storage Manager storage, but leaving the original file intact on the local file system.

premigration percentage

A space management setting that controls whether the next eligible candidates in a file system are premigrated following threshold or demand migration.

primary storage pool

A named set of volumes that the server uses to store backup versions of files, archive copies of files, and files migrated from client nodes. See also *destination* and *copy storage pool*.

privilege class

A level of authority that is granted to an administrator. The privilege class determines which administrative tasks the administrator can perform. See also *node privilege class*, *operator privilege class*, *policy privilege class*, *storage privilege class*, and *system privilege class*.

profile

A named group of configuration information that can be distributed from a configuration manager when a managed server subscribes. Configuration information can include registered administrator IDs, policies, client schedules, client option sets, administrative schedules, storage manager command scripts, server definitions, and server group definitions. See also *configuration manager* and *managed server*.

Q

quota (1) For HSM on AIX, UNIX, or Linux systems, the limit (in megabytes) on the amount of data that can be migrated and premigrated from a file system to server storage.

(2) For HSM on Windows systems, a user-defined limit to the space that is occupied by recalled files.

R

randomization

The process of distributing schedule start

times for different clients within a specified percentage of the schedule's startup window.

raw logical volume

A portion of a physical volume that is comprised of unallocated blocks and has no journaled file system (JFS) definition. A logical volume is read/write accessible only through low-level I/O functions.

read-without-recall recall mode

A mode that causes hierarchical storage management (HSM) to read a migrated file from storage without storing it back on the local file system. The last piece of information read from the file is stored in a buffer in memory on the local file system. Contrast with *normal recall mode* and *migrate-on-close recall mode*.

rebind

To associate a backed-up file with a new management class name. For example, rebinding occurs when the management class associated with a file is deleted. See also *bind*.

recall In Tivoli Storage Manager, to copy a migrated file from server storage back to its originating file system using the space management client. See also *transparent recall*, *selective recall*, and *recall mode*.

recall mode

A mode that is assigned to a migrated file with the `dsmattr` command that determines how the file is processed when it is recalled. It determines whether the file is stored on the local file system, is migrated back to Tivoli Storage Manager storage when it is closed, or is read from Tivoli Storage Manager storage without storing it on the local file system.

receiver

A server repository that contains a log of server and client messages as events. For example, a receiver can be a file exit, a user exit, or the Tivoli Storage Manager server console and activity log. See also *event*.

reclamation

The process of consolidating the remaining data from many sequential-access volumes onto fewer, new sequential-access volumes.

reclamation threshold

The percentage of space that a sequential-access media volume must have before the server can reclaim the volume. Space becomes reclaimable when files are expired or are deleted.

reconciliation

The process of synchronizing a file system with the Tivoli Storage Manager server, and then removing old and obsolete objects from the Tivoli Storage Manager server.

recovery log

A log of updates that are about to be written to the database. The log can be used to recover from system and media failures. The recovery log consists of the active log (including the log mirror) and archive logs.

register

To define a client node or administrator ID that can access the server.

registry

A repository that contains access and configuration information for users, systems, and software.

resident file

On a Windows system, a complete file on a local file system that might also be a migrated file because a migrated copy can exist in Tivoli Storage Manager storage. On a UNIX or Linux system, a complete file on a local file system that has not been migrated or premigrated, or that has been recalled from Tivoli Storage Manager storage and modified. Contrast with *stub file* and *premigrated file*. See *migrated file*.

restore

To copy information from its backup location to the active storage location for use. For example, to copy information from server storage to a client workstation.

retention

The amount of time, in days, that inactive backed-up or archived files are kept in the storage pool before they are deleted. Copy group attributes and default retention grace periods for the domain define retention.

retrieve

To copy archived information from the

storage pool to the workstation for use. The retrieve operation does not affect the archive version in the storage pool.

roll back

To remove changes that were made to database files since the last commit point.

root user

A system user who operates without restrictions. A root user has the special rights and privileges needed to perform administrative tasks.

S

SAN See *storage area network*.

schedule

A database record that describes client operations or administrative commands to be processed. See *administrative command schedule* and *client schedule*.

scheduling mode

The type of scheduling operation for the server and client node that supports two scheduling modes: client-polling and server-prompted.

scratch volume

A labeled volume that is either blank or contains no valid data, that is not defined, and that is available for use.

script A series of commands, combined in a file, that carry out a particular function when the file is run. Scripts are interpreted as they are run. Contrast with *Tivoli Storage Manager command script*.

Secure Sockets Layer (SSL)

A security protocol that provides communication privacy. With SSL, client/server applications can communicate in a way that is designed to prevent eavesdropping, tampering, and message forgery.

selective backup

The process of backing up certain files or directories from a client domain. The files that are backed up are those that are not excluded in the include-exclude list. The files must meet the requirement for serialization in the backup copy group of the management class that is assigned to each file. Contrast with *incremental backup*.

selective migration

The process of copying user-selected files

- from a local file system to Tivoli Storage Manager storage and replacing the files with stub files on the local file system. Contrast with *threshold migration* and *demand migration*.
- selective recall**
The process of copying user-selected files from Tivoli Storage Manager storage to a local file system. Contrast with *transparent recall*.
- serialization**
The process of handling files that are modified during backup or archive processing. See *dynamic serialization*, *static serialization*, *shared static serialization*, and *shared dynamic serialization*.
- server** A software program or a computer that provides services to other software programs or other computers.
- server options file**
A file that contains settings that control various server operations. These settings affect such things as communications, devices, and performance.
- server-prompted scheduling mode**
A client/server communication technique where the server contacts the client node when tasks must be done. Contrast with *client-polling scheduling mode*.
- server storage**
The primary, copy, and active-data storage pools that are used by the server to store user files such as backup versions, archive copies, and files migrated from space manager client nodes (space-managed files). See also *active-data pool*, *primary storage pool*, *copy storage pool*, *storage pool volume*, and *volume*.
- session**
A logical or virtual connection between two stations, software programs, or devices on a network that allows the two elements to communicate and exchange data.
- session resource usage**
The amount of wait time, processor time, and space that is used or retrieved during a client session.
- shared dynamic serialization**
A value for serialization that specifies that a file must not be backed up or archived if it is being modified during the operation. Tivoli Storage Manager retries the backup or archive operation a number of times; if the file is being modified during each attempt, Tivoli Storage Manager will back up or archive the file on its last try. See also *serialization*. Contrast with *dynamic serialization*, *shared static serialization*, and *static serialization*.
- shared library**
A library device that is used by multiple storage manager servers.
- shared static serialization**
A copy-group serialization value that specifies that a file must not be modified during a backup or archive operation. Tivoli Storage Manager attempts to retry the operation a number of times. If the file is in use during each attempt, the file is not backed up or archived. See also *serialization*. Contrast with *dynamic serialization*, *shared dynamic serialization*, and *static serialization*.
- snapshot**
An image backup type that consists of a point-in-time view of a volume.
- space-managed file**
A file that is migrated from a client node by the space manager client. The space manager client recalls the file to the client node on demand.
- space management**
The process of keeping sufficient free storage space available on a local file system for new data by migrating files to server storage. Synonymous with *hierarchical storage management*.
- space manager client**
A program that runs on a UNIX or Linux system to manage free space on the local file system by migrating files to server storage. The program can recall the files either automatically or selectively. Also called *hierarchical storage management (HSM) client*.
- space monitor daemon**
A daemon that checks space usage on all file systems for which space management is active, and automatically starts threshold migration when space usage on a file system equals or exceeds its high threshold.

sparse file

A file that is created with a length greater than the data it contains, leaving empty spaces for the future addition of data.

special file

On AIX, UNIX, or Linux systems, a file that defines devices for the system, or temporary files that are created by processes. There are three basic types of special files: first-in, first-out (FIFO); block; and character.

SSL See *Secure Sockets Layer*.

stabilized file space

A file space that exists on the server but not on the client.

stanza A group of lines in a file that together have a common function or define a part of the system. Each stanza is identified by a name that occurs in the first line of the stanza. Depending on the type of file, a stanza is ended by the next occurrence of a stanza name in the file, or by an explicit end-of-stanza marker. A stanza can also be ended by the end of the file.

startup window

A time period during which a schedule must be initiated.

static serialization

A copy-group serialization value that specifies that a file must not be modified during a backup or archive operation. If the file is in use during the first attempt, the storage manager cannot back up or archive the file. See also *serialization*. Contrast with *dynamic serialization*, *shared dynamic serialization*, and *shared static serialization*.

storage agent

A program that enables the backup and restoration of client data directly to and from storage attached to a storage area network (SAN).

storage area network (SAN)

A dedicated storage network that is tailored to a specific environment, combining servers, systems, storage products, networking products, software, and services.

storage hierarchy

(1) A logical order of primary storage pools, as defined by an administrator. The

order is typically based on the speed and capacity of the devices that the storage pools use. The storage hierarchy is defined by identifying the next storage pool in a storage pool definition. See also *storage pool*.

(2) An arrangement of storage devices with different speeds and capacities. The levels of the storage hierarchy include: main storage, such as memory and direct-access storage device (DASD) cache; primary storage (DASD containing user-accessible data); migration level 1 (DASD containing data in a space-saving format); and migration level 2 (tape cartridges containing data in a space-saving format).

storage pool

A named set of storage volumes that are the destination that is used to store client data. A storage pool contains backup versions, archive copies, and files that are migrated from space manager client nodes. A primary storage pool is backed up to a copy storage pool. See also *primary storage pool*, *copy storage pool*, and *active-data pool*.

storage pool volume

A volume that has been assigned to a storage pool. See also *volume*, *active-data pool*, *copy storage pool*, and *primary storage pool*.

storage privilege class

A privilege class that gives an administrator the authority to control how storage resources for the server are allocated and used, such as monitoring the database, the recovery log, and server storage. See also *privilege class*.

stub

A shortcut on the Windows file system that is generated by the hierarchical storage management (HSM) client for a migrated file that allows transparent user access. A stub is the sparse file representation of a migrated file, with a reparse point attached.

stub file

A file that replaces the original file on a local file system when the file is migrated to storage. A stub file contains the information that is necessary to recall a migrated file from Tivoli Storage Manager storage. It also contains additional

information that can be used to eliminate the need to recall a migrated file.

stub file size

The size of a file that replaces the original file on a local file system when the file is migrated to Tivoli Storage Manager storage. The size that is specified for stub files determines how much leader data can be stored in the stub file. The default for stub file size is the block size defined for a file system minus 1 byte.

subscription

In a Tivoli environment, the process of identifying the subscribers that the profiles are distributed to. For Tivoli Storage Manager, a subscription is the process by which a managed server receives configuration information associated with a particular profile on a configuration manager. See also *managed server*, *configuration manager*, and *profile*.

Systems Network Architecture (SNA)

The description of the logical structure, formats, protocols, and operational sequences for transmitting information through and controlling the configuration and operation of networks.

system privilege class

A privilege class that gives an administrator the authority to issue all server commands. See also *privilege class*.

T

tape library

A set of equipment and facilities that support an installation's tape environment. The tape library can include tape storage racks, mechanisms for automatic tape mounting, a set of tape drives, and a set of related tape volumes mounted on those drives.

tape volume prefix

The high-level-qualifier of the file name or the data set name in the standard tape label.

target node

A client node for which other client nodes (called agent nodes) have been granted proxy authority. The proxy authority allows the agent nodes to perform operations such as backup and restore on behalf of the target node, which owns the data.

TCA See *trusted communications agent*.

TCP/IP

See *Transmission Control Protocol/Internet Protocol*.

threshold migration

The process of moving files from a local file system to Tivoli Storage Manager storage based on the high and low thresholds that are defined for the file system. Contrast with *demand migration*, *selective migration*, and *migration job*.

throughput

In storage management, the total bytes in the workload, excluding overhead, that are backed up or restored, divided by elapsed time.

timeout

A time interval that is allotted for an event to occur or complete before operation is interrupted.

timestamp control mode

A mode that determines whether commands preserve the access time for a file or set it to the current time.

Tivoli Storage Manager command script

A sequence of Tivoli Storage Manager administrative commands that are stored in the database of the Tivoli Storage Manager server. The script can run from any interface to the server. The script can include substitution for command parameters and conditional logic.

tombstone object

A small subset of attributes of a deleted object. The tombstone object is retained for a specified period, and at the end of the specified period, the tombstone object is permanently deleted.

Transmission Control Protocol/Internet Protocol (TCP/IP)

An industry-standard, nonproprietary set of communication protocols that provides reliable end-to-end connections between applications over interconnected networks of different types.

transparent recall

The process that is used to automatically recall a file to a workstation or file server when the file is accessed. See also *recall mode*. Contrast with *selective recall*.

trusted communications agent (TCA)

A program that handles the sign-on password protocol when clients use password generation.

U

UCS-2 A 2-byte (16-bit) encoding scheme based on ISO/IEC specification 10646-1. UCS-2 defines three levels of implementation: Level 1-No combining of encoded elements allowed; Level 2-Combining of encoded elements is allowed only for Thai, Indic, Hebrew, and Arabic; Level 3-Any combination of encoded elements are allowed.

UNC See *Universal Naming Convention name*.

Unicode

A character encoding standard that supports the interchange, processing, and display of text that is written in the common languages around the world, plus some classical and historical texts. The Unicode standard has a 16-bit character set defined by ISO 10646.

Unicode-enabled file space

Unicode file space names provide support for multilingual workstations without regard for the current locale.

Unicode transformation format 8

Unicode Transformation Format (UTF), 8-bit encoding form, which is designed for ease of use with existing ASCII-based systems. The CCSID value for data in UTF-8 format is 1208.

Universal Naming Convention (UNC) name

A name that is used to access a drive or directory containing files shared across a network. The UNC name includes the system name and a SharePoint name that represents the shared drive or directory.

Universally Unique Identifier (UUID)

The 128-bit numerical identifier that is used to ensure that two components do not have the same identifier.

UTF-8 See *Unicode transformation format 8*.

UUID See *Universally Unique Identifier*.

V**validate**

To check a policy set for conditions that can cause problems if that policy set becomes the active policy set. For

example, the validation process checks whether the policy set contains a default management class.

version

A backup copy of a file stored in server storage. The most recent backup copy of a file is the active version. Earlier copies of the same file are inactive versions. The number of versions retained by the server is determined by the copy group attributes in the management class.

virtual file space

A representation of a directory on a network-attached storage (NAS) file system as a path to that directory.

virtual volume

An archive file on a target server that represents a sequential media volume to a source server.

volume

A discrete unit of storage on disk, tape or other data recording medium that supports some form of identifier and parameter list, such as a volume label or input/output control. See also *scratch volume*, and *storage pool volume*.

volume history file

A file that contains information about volumes that have been used by the server for database backups and for export of administrator, node, policy, or server data. The file also has information about sequential-access storage pool volumes that have been added, reused, or deleted. The information is a copy of volume information that is recorded in the server database.

Volume Shadow Copy Service

A set of Microsoft application-programming interfaces (APIs) that you can use to create shadow copy backups of volumes, exact copies of files, including all open files, and so on.

VSS See *Volume Shadow Copy Service*.

VSS Backup

A backup operation that uses Microsoft Volume Shadow Copy Service (VSS) technology. The backup operation produces an online snapshot (point-in-time consistent copy) of Exchange data. This copy can be stored

on local shadow volumes or on Tivoli Storage Manager server storage.

VSS Fast Restore

A function that uses a Microsoft Volume Shadow Copy Service (VSS) software provider to restore VSS Backups (IBM Data Protection for Exchange database files and log files) that reside on local shadow volumes.

VSS Instant Restore

A volume-level hardware-assisted Microsoft Volume Shadow Copy Service (VSS) function where target volumes that contain the snapshot are copied back to the original source volumes.

VSS offloaded backup

A backup operation that uses a Microsoft Volume Shadow Copy Service (VSS) hardware provider (installed on an alternate system) to move IBM Data Protection for Exchange data to the Tivoli Storage Manager server. This type of backup operation shifts the backup load from the production system to another system.

VSS Restore

A function that uses a Microsoft Volume Shadow Copy Service (VSS) software provider to restore VSS Backups (IBM Data Protection for Exchange database files and log files) that reside on Tivoli Storage Manager server storage to their original location.

W

wildcard character

A special character such as an asterisk (*) or a question mark (?) that can be used to represent one or more characters. Any character or set of characters can replace the wildcard character.

workstation

A configuration of input/output equipment at which an operator works. A workstation is a terminal or microcomputer at which a user can run applications and that is usually connected to a mainframe or a network.

worldwide name

A 64-bit, unsigned name identifier that is unique.

workload partition (WPAR)

A partition within a single operating system instance.

Index

Special characters

.SpaceMan directory 51, 137
.SpaceMan/orphan.stubs 90

A

accessibility features 139
activating space management 116
add from HSM GUI 48
adding space management to 49
adding space management to a file system 112
adding to nested file systems 49
AIX GPFS 8
 installing 7
 installing HSM 8, 12
 requirements 8
 upgrading HSM client 9
AIX JFS2
 hardware requirements 11
 installing 10
 requirements 11
 software requirements 11
archiving migrated files 68
automatic migration 2, 62
 file prioritization 60
 threshold migration 107
automatic reconciliation 79
automatically 62
automignonuse 33
automigration 63

B

backing up
 premigrated files 69
backing up files 67
 after migrating 69
 before migrating 69
backup and migration 72, 73
backup-archive client 1
Bourne and Korn shell 37
Bourne shell variables 37

C

C shell 38
C shell variables 38
candidate selection
 migration 61
candidates file 137
candidatesinterval option 30, 89
checkfororphans option 30, 80, 90
checkthresholds option 30, 90
child recall daemon 83
CIFS 49
client system options
 maxrecalldaemons 94

closed registration
 HSM UNIX client 22
clusters on GPFS
 adding space management 42
command
 dsmautomig 107
 dsmmighelp 102
 dsmmigquery -options 28
 dsmmigundelete 70
command line help
 displaying 102
commands 70
descriptions 103
dmkilld 105
dsmattr 76, 78, 105
dsmautomig 80, 94, 107
dsmc 88
dsmdf 59, 103, 108
dsmdu 59, 103, 109
dsmls 103, 110, 111, 128
dsmmigfs 57
dsmmigfs 51, 56, 76, 97, 103, 117, 120
 adding or updating 112
 adding space management 48
 deactivating space management 116
 globally deactivating space management 119
 reactivating space management 116
 removing space management 116
 takeover 121
dsmmigfs command
 globally deactivating space management 119
 globally reactivating space management 119
dsmmigfs rollback 120
dsmmigfs SDRReset 118
dsmmighelp 122
dsmmigquery 34, 122, 128
dsmmigrate 51, 63, 124
dsmmigundelete 103, 126
dsmmonitor 83, 127
dsmq 85, 127, 131
dsmrecall 128
dsmrecalld 83, 130, 134
dsmreconcile 80, 81, 90, 130
dsmrm 85, 131
dsmrootd 85, 132
dsmscoutd 132
dsmscoutd command 134
dsmsetpw 133
dsmwatchd 134
 using 101
compressalways option 91
compression 21
compression option 30, 91
creating stub files for premigrated files 126

customer support
 contact xv

D

daemon
 dsmmonitor 107
 dsmscout 53
 dsmscoutd 61
daemons
 dsmmigrate 120
 dsmmonitor 127, 130
 dsmrecall 120
 dsmscoutd 130
 dsmrootd 132
 dsmscoutd 132
 dsmwatchd 120, 134
 dsmwatchd daemon 47
 overview 83
 scout 84
 stopping 85
deactivating
 space management 55
deactivating space management 116
 globally 119
default size 55
defaultserver option 30, 91
defined 3
definitions 145
deleted stub files 70
deleted stub files, restoring 70, 126
demand migration 2, 60
disable failover 118
disaster recovery 74
disk space requirements 19
disk, restoring 74
dismmigfs
 command
 reactivate space management 56
displaying 28
 space management information 109
 space management settings 117
 space usage information 109
displaying online command help 102
displaying option information 28
DMAPI 47
 mmchfs command 8
dmkilld command 105
dms.sys
 reconcileinterval 83
dmsreconcile command 90
DSM_CONFIG 37
DSM_CONFIG environment variable 37
DSM_DIR environment variable 37
DSM_LOG environment variable 37
dsm.opt 5, 89
 editing 28, 32
dsm.sys 5, 61, 89
 editing 28
 maxcandprocs 61
 maxrecalldaemons 83

- dsm.sys (*continued*)
 - minrecalldaemons 83
 - options 29
- dsmattr command 76, 78, 105
 - file list 92
- dsmautomig command 62, 94, 107
- dsmc command 88
- dsmdf command 103, 108
- dsmdu command 103, 109
- dsmerror.log 37
- dsmls command 103, 110, 111, 128
 - file list 92
- dsmmigfs
 - removing space management 57
- dsmmigfs 63
 - command 121
 - deactivating space management 116
 - GPFS failover 118
 - query 117
 - reactivating space management 116
- dsmmigfs command 48, 51, 76, 97, 103, 120
 - adding space management 112
 - updating space management settings 112
- dsmmigfs DISABLEFailover 118
- dsmmigfs ENABLEFailover 118
- dsmmigfs SDRReset 118
- dsmmighelp command 102, 122
- dsmmigquery
 - shared options 123
- dsmmigquery -options 28
- dsmmigquery command 34, 122, 128
- dsmmigrate 63
- dsmmigrate command 51, 63, 124
 - file list 92
- dsmmigundelete 70
- dsmmigundelete command 70, 103, 126
- dsmmonitord command 84, 127
 - space monitor daemon 83
- dsmmonitord daemon 107, 130
- DSMNodeset file 14
- dsmq command 85, 127, 131
- dsmrecaid daemon 7
- dsmrecall command 128
 - file list 92
- dsmrecalld command 83, 130, 134
- dsmreconcile
 - command 81
 - reconcilinterval 3
- dsmreconcile command 80, 130
- dsmrm command 85, 131
- dsmrootd command 85, 132
- dsmsched.log 88
- dsmscoutd
 - daemon 61
- dsmscoutd command 132, 134
- dsmscoutd daemon 30, 89
- DSMSDRVersion file 14
- dsmsetpw command 133
- dsmwatchd command 134

E

- editing
 - dsm.opt 28
 - dsm.sys 28

- education
 - see Tivoli technical training xiii
- enable failover 118
- enablelanfree 38
- encrypted file system 112
- entire system 74
- environment 43
- environment variables
 - DSM_CONFIG 37
 - DSM_DIR 37
 - DSM_LOG 37
- errorprog option 30, 92
- exclude
 - files (LAN-free) 39
- exclude list
 - edit 36
- exclude options 35
- exclude.compression 91
- excluding files
 - from backup 35
 - from migration 35
 - from space management 35
- expiration 3
 - file 68
- exporting file systems 49
- exporting files using
 - NFS 49
- external storage pool 63

F

- failover
 - watch daemon 84
- FastConnect 49
- file
 - hsmfsconfig.xml 51
- file list 111
- file size
 - minimum partial 55
- file systems 47, 72
 - reconcile 3
 - reconciling 79
- filelist option 30, 92
- files
 - backing up 67
 - migrating 67
 - restoring 67
- fixes, obtaining xiv
- for AIX GPFS 8
- for Linux x86/x86_64 GPFS 15

G

- globally deactivating space management 119
- globally reactivating space management 119
- glossary 145
- GPFS 63
- GPFS policy 63
- GPFS rule 63

H

- HACMP
 - limitations 40

- HACMP for AIX JFS2
 - automatic migration 39
 - enable failover 39
 - to source nodes in cluster 39
- hardware requirements 8, 19
 - for AIX JFS2 11
- hardware, disk space, memory requirements 19
- help 102
- high threshold 52
 - setting 112
- HP Itanium 64
 - installing 13
- HP-UX IA64
 - requirements 13
- HSM
 - installing on Solaris 19
- hsm agent
 - verify running 24
- HSM client
 - installing 5
 - open registration on TSM 21
- HSM client overview 1
- HSM GUI 23
- HSM Solaris client 19
 - requirements 19
- HSM UNIX client
 - closed registration on TSM 22
- hsmagent.opt 44
- hsmfsconfig.xml 47
- hsmfsconfig.xml file 51

I

- IBM Software Support
 - submitting a problem xvi
- IBM Support Assistant xiv
- ILM 63
- includexcl option 30, 93
 - dsm.sys 35
- include
 - files (LAN-free) 39
- include list
 - edit 36
- include options 35
- include-exclude
 - filed 27
- include-exclude files 5
- include.compression 91
- including files
 - for backup 35
 - for migration 35
- Information Lifecycle Management 63
- inodes 108
- installation requirements
 - for AIX GPFS 7
 - for HP Itanium 64 13
 - for Solaris 18
- installing 14, 18
 - HP-UX IA64 13
 - Linux x86/x86_64 GPFS 16
 - prerequisites 5
- installing HSM 5, 14, 18
 - AIX GPFS 8, 12
 - AIX JFS2 10
 - on AIX GPFS 7
 - on HP Itanium 64 13

Internet, searching for problem resolution xiii, xiv

J

JFS2 EFS 112

K

knowledge bases, searching xiii
Korn shell variables 37

L

LAN-free 35, 38
 prerequisites 38
LAN-free data transfer 38
lanfreecommmethod 38
lanfreetcpport 38
Linux GPFS
 requirements 15
 upgrading HSM client 16
Linux x86/x86_64 GPFS 14, 15
list of files
 displaying 122
local zones
 Solaris 20
log files
 dsmerror.log 37
logdir directory 137
logs
 dsmsched 88
low threshold 52
 setting 112

M

manage recovery 118
management 47
management class
 assign to files 33
 default 33
 displaying 34
manually 62
maxcandidates
 setting 112
maxcandprocs option 30, 94
maxmigrators option 31, 94
maxrecalldaemons
 option 83
maxrecalldaemons option 31, 94
maxreconcileproc option 31, 95
maxthresholdproc option 31, 95
memory requirements 19
migdestination 33
migfileexpiration
 option 79
migfileexpiration option 31, 68, 80, 96, 99
migrated
 orphans 79
migrated file
 setting a recall mode 105
migrated files
 recalling 75

migrateserver option 31, 91, 96
migrating
 archived files 68
 demand 2
 overview 2
 recalling files 3
 selective 2
 threshold 2
migrating files 67
 archiving 68
 automatic 60
 automatically 107
 demand 60
 file list 92
 recording information 137
 selectively 124
 threshold 60
 threshold migration 62
migrating from AIX JFS 12
migration 52
 automatic 59
 automatic migration 2
 candidate selection 61
 overview 59
 premigration 2
 selective 59
 selective migration 124
 threshold migration 107
migration candidates list 80, 130, 137
 displaying 122
migrequiresbkup 33
minimum file size 52
minimum file size for migration 52
minimum size 55
minmigfilesize option 31, 97
minpartialrecallsz option 76
minrecalldaemons
 option 83
minrecalldaemons option 31, 97
mmchpolicy 63

N

nested file system
 mounting parent first 49
NFS servers 49
normal recall
 mode 76
 setting for a migrated file 105

O

obsolete files
 reconciliation 3
online help
 displaying for commands 102, 122
open registration
 HSM client 21
operating system 74
optional setup 32
optionformat option 31, 98
options 28, 38, 89
 candidatesinterval 29, 84, 89
 checkfororphans 29, 80, 81, 90
 checkthresholds 29, 90
 compressalways 91

options (*continued*)
 compression 29, 91
 defaultserver 29, 91
 enablelanfree 38
 errorprog 29, 81, 92
 exclude 35
 filelist 92
 for scheduling services 87
 inlexcl 93
 include 35
 kernelmessages 29
 lanfreecommmethod 38
 lanfreetcpport 38
 maxcandprocs 94
 maxmigrators 29, 94
 maxrecalldaemons 29, 83, 94
 maxreconcileproc 29, 95
 maxthresholdproc 29, 95
 migfileexpiration 29, 68, 79, 80, 96, 99
 migrateserver 29, 91, 96
 minmigfilesize 97
 minrecalldaemons 29, 31, 83, 97
 optionformat 98
 passwordaccess 22
 reconcileinterval 29, 79, 87, 98
 reconciliation
 tasks 79
 restoremigrate 68
 restoremigstate 29, 99
 restoreMigState 80
 rollback 120
 schedlogname 88
 selvelastaccessdate 59
 setting space management 27
 standard format 101
ordered recall list
 displaying 122
orphan stub files 81
orphan stubs 80
 identifying 137
orphan.stubs file 79, 137
orphans 79
overview 5

P

packages
 AIX JFS2 11
parent
 mounting before nested file systems 49
partial file
 minimum size 55
partial file recall 3, 76
partial recall mode
 setting for a migrated file 105
password 133
 encryption 22
 generate 22
 passwordaccess 22
 TSM 22
policies
 displaying 33, 34
premigrated files
 database 137
 orphaned 79

- premigrating
 - files 62, 67, 69
- premigrating files
 - archiving 68
- premigration 2, 62, 63
- premigration percentage 53, 62, 63
- premigration setting 112
- premgrdb.dir file 137
- premgrdb.pag file 137
- prerequisites 38
 - installing 5
 - LAN-free 38
- problem determination
 - describing problem for IBM Software Support xvi
 - determining business impact for IBM Software Support xv
 - submitting a problem to IBM Software xvi
- publications
 - download xi
 - order xi
 - search xi
 - Tivoli Storage Manager xi

Q

- quota 54
 - setting 112

R

- reactivating space management 56, 116
 - globally 119
- read-only recall 76, 77
- recall
 - daemon 83
 - HACMP 39
- recall daemon 31
 - dmkill command 105
 - starting 130
- recall ID 127
- recall mode
 - dsmattr command 78
 - HSM determines 78
 - normal 105
 - setting for a migrated file 105
- recall process
 - displaying status of 127
- recalling files 76, 78
 - canceling 131
 - file list 92
 - migrated files 3
 - normal 76
 - recording information 137
 - removing from queue 131
 - selective recall 75
 - selectively 76, 128
 - streaming 77
 - transparent recall 75
- recallmode
 - option 76
- recommendations
 - not adding space management 48
- reconcile file systems 3

- reconcileinterval
 - dms.sys 83
 - option 79
- reconcileinterval option 31, 87, 98
- reconciliation
 - automatic 79
- reconciling
 - manual 79
 - manually 80, 130
- reconciling file systems 130
- recreating deleted stub files 70, 126
- removing
 - space management 116
- requirements 19
 - AIX GPFS 8
 - AIX JFS2 11
 - HP-UX IA64 13
 - installing 5
 - Linux GPFS 15
- reset in the SDR 118
- resolving 80
- Responsiveness Service 84
- restarting space management daemons 120
- restore migrated files 70
- restoremigrate 68
- restoremigstate
 - command 81
- restoremigstate option 31, 99
- restoreMigState option 80
- restoring 67, 70, 72, 74
 - a disk 74
 - deleted stub files 126
 - premigrated files 69
- restoring files 67
- retrieving
 - archived files 68
- rollback option 120
- root daemon 85
 - starting 132
- root user tasks 127
 - creating stub files for premigrated files 126
 - recreating deleted stub files 126
 - starting a root daemon 132
 - starting a scout daemon 132

S

- SAMBA 49
- sample script 40
- SAN
 - LAN-free 38
- schelogname option 88
- scheduled services
 - displaying 88
 - displaying completed services 88
- scheduler
 - starting 87
- scheduling
 - backup-archive services 87
 - space management services 87
- scheduling services
 - options 87
- scout daemon 84
 - starting 132

- selected files
 - migrating 124
 - recalling 128
- selecting HSM options 27
- selective migration 2, 63
 - performing 124
- selective recall 3, 76, 128
- server is different 73
- server is the same 72
- services
 - scheduling 87
- setting
 - quota 54
- setting environment variables 37
- setting space management options 27
- setting thresholds for 52
- setting up 27
- settings
 - minimum partial file size 55
 - minimum streaming file size 55
 - premigration percentage 53
 - updating from HSM GUI 50
- setup tasks 32
- shell scripts 1
- software requirements 8, 15
 - for AIX JFS2 11
 - HSM Solaris client 19
- Software Support
 - contact xv
 - describing problem for IBM Software Support xvi
 - determining business impact for IBM Software Support xv
- Solaris
 - installing HSM 19
 - requirements 19
 - uninstalling HSM client 21
 - upgrading HSM client 20
- Solaris local zone support 20
- space management 27, 48, 49
 - adding to a file system 112
 - agent 44
 - clusters on GPFS 42
 - files 137
 - candidates file 137
 - logdir directory 137
 - orphan.stubs file 137
 - premgrdb.dir file 137
 - premgrdb.pag file 137
 - status file 137
 - installing 5
 - maintaining 112
 - removing 57, 116
 - settings 51
 - updating a file system 112
- space management settings 112
 - displaying 117
 - high threshold 112
 - low threshold 112
 - maxcandidates 112
 - premigration setting 112
 - quota 112
 - stub file size 112
- space monitor daemon 30, 83, 90
 - starting 127
- space usage information
 - displaying 109

- space usage, displaying 103, 108
- SpaceMan directory
 - contents 135
 - setting low threshold 52
- spacemgtechnique 33
- standard option format 101
- start HACMP 40
- starting
 - client scheduler 87
 - space management daemons 120
- status file 137
- stop HACMP 40
- stopping space management
 - daemons 85, 120
- storage pool 1, 83, 84
 - dsmautomig command 107
- storage pools 62
- stream file size 55
- streaming recall 3
- streaming recall mode 77
- stub file size
 - setting 112
- stub files
 - default size by platform 54
 - leader data 54
 - maximum size 54
 - minimum size 54
 - orphan 81
 - recreating 126
 - recreating deleted 103
 - restoring 70
 - setting the size 54
- summary of changes V6.1 xix
- support information xiii
- synchronizing client-server files 79
 - manual reconciliation 130

T

- takeover filespec 121
- tasks 32
- threshold migration 2, 60
 - start manually 62
 - starting manually 107
- threshold rule 63
- thresholds
 - migration 52
- Tivoli Storage Manager password 133
- Tivoli technical training xiii
- training, Tivoli technical xiii
- transfer HSM management of file
 - system 120, 121
- transparent recall 3, 75
- truncate
 - migrated files 3

U

- uninstalling
 - backup-archive (Solaris) 20
 - from AIX GPFS 10
 - HP-UX IA HSM client 13
 - HSM client from Solaris 21
- uninstalling from Linux 17
- update files in SDR 118

- updating space management
 - settings 112
- upgrading
 - HP-UX IA HSM client 13
 - Solaris HSM client 20
- upgrading from AIX JFS 12
- upgrading HSM client
 - AIX GPFS 9
 - Linux GPFS 16
- using commands 101

V

- variables 37, 38
- variables, environment 37
- verifying
 - hsm agent running 24

W

- watch daemon 47, 84
- workstation
 - registering with a server 21
- WPAR 50



Program Number: 5608-E01, 5608-E02, 5608-E07, 5608-E12

Printed in USA

SC23-9794-01

