

Tivoli Storage Manager for Space Management for UNIX
and Linux
Version 6.3

User's Guide



Tivoli Storage Manager for Space Management for UNIX
and Linux
Version 6.3

User's Guide



Note:

Before using this information and the product it supports, read the information in "Notices" on page 177.

This edition applies to Version 6.3 of IBM Tivoli Storage Manager for Space Management (product number 5608-E12), and to all subsequent releases and modifications until otherwise indicated in new editions or technical newsletters. This edition replaces SC23-9794-02.

© **Copyright IBM Corporation 1995, 2011.**

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	xii
Support information	xiv
Getting technical training	xiv
Searching knowledge bases	xv
Contacting IBM Software Support	xvii
Conventions used in this publication	xviii

Tivoli Storage Manager for Space Management version 6.3 updates	xix
--	------------

Chapter 1. HSM client for UNIX and Linux overview	1
--	----------

Migrating files overview	2
Recalling migrated files overview	3
Reconciling file systems overview	4

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	9
Installing the AIX HSM client	9
Upgrading the AIX GPFS HSM client	10
Uninstalling the AIX GPFS HSM client	11
AIX JFS2 HSM client installation overview	11
AIX JFS2 HSM client environment requirements	12
Installing the AIX HSM client	12
Linux x86_64 GPFS HSM client installation overview	13
Linux x86_64 GPFS HSM client environment requirements	15
Installing the Linux x86_64 GPFS client	15
Upgrading the Linux x86_64 GPFS HSM client	16
Uninstalling the Linux x86_64 GPFS HSM client	17
Registering your workstation with a Tivoli Storage Manager server	18
Open registration	18
Closed registration	19
Managing your password	19
The HSM GUI	20
Verifying the HSM agent is running	20
Using the HSM GUI	21

Chapter 3. Configuring the HSM client	23
--	-----------

Displaying HSM client options	24
---	----

Editing the options files using the backup-archive client	24
Editing the options file dsm.sys	25
Editing the options file dsm.opt	28
Optional setup features	29
Assigning management classes to files	29
Displaying management class information	30
Modifying the include-exclude options file	31
Setting environment variables	33
Setting up LAN-free data transfer for HSM.	34
Setting up space management for HACMP on AIX JFS2	35
Space management for AIX GPFS and Linux x86_64 GPFS clusters	38
Configuring the HSM space management agent	40
Setting HSM logging	41

Chapter 4. Adding and configuring space management for a file system	43
---	-----------

Adding space management to file systems	44
Adding space management to nested file systems	45
Mounting the parent file system before nested file system.	45
Adding space management to an exported file system	45
Removing or reconfiguring a space-managing node in a GPFS cluster	46
Adding space management to WPARS on AIX 6.1 and 7.1	46
Updating settings from the command line	47
Updating settings from the HSM GUI	47
Settings that control the space usage of a file system	47
Minimum migration file size	49
Migration threshold percentages	49
Premigration percentage	50
Quotas	51
Stub file size	52
Minimum stream file size.	52
Minimum partial file size.	53
Maximum number of files	53

Managing a file system with multiple Tivoli Storage Manager servers.	53
Enabling a file system to be managed by multiple Tivoli Storage Manager servers	54
Adding space management to a file system in a multiple-server environment.	55
Removing a Tivoli Storage Manager server from a multiple-server environment	56
Limitations for multiple-server environments	57
Replicating nodes in a multiple-server environment	57
Deactivating space management	58
Reactivating space management	59
Removing space management	60

Chapter 5. Migrating files 61
 Automatic file migration 62
 Candidate selection for automatic migration 63
 File premigration 64
 Manually starting threshold migration 65
 Migrating selectively using the dsmmigrate
 command 65
 | Migrating a single file system to two or more Tivoli
 | Storage Manager servers 66

**Chapter 6. Backing up and restoring
 file systems 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 file system in an environment that is
 | managed by multiple Tivoli Storage Manager
 | servers 74
 | Restoring a disk 74
 | Restoring your operating system and file systems 75

Chapter 7. Recalling migrated files. . . . 77
 Transparent recall processing 77
 Selective recall processing 77
 Normal recall mode 78
 Partial file recall mode. 78
 Streaming recall mode. 79
 How HSM determines which recall mode to use . . 80
 Setting the recall mode using the dsmmatr command 80
 | Optimized tape recall processing 80
 | List files for optimized tape recalls 81
 | Recalling files with optimization for tape 84

Chapter 8. Reconciling file systems . . . 85
 Automatic reconciliation 85
 Manually reconciling file systems 86
 Resolving orphaned stub files 87
 | Reconciliation based on a GPFS policy 87

**Chapter 9. Space management
 daemons 89**
 The space monitor daemon 89
 The recall daemon 90
 The scout daemon 90
 The watch daemon 91
 The root daemon 91
 Stopping the space management daemons 92

Chapter 10. Scheduling services. 93
 Scheduling options 93
 Starting the client scheduler 93
 Displaying scheduled services information 94

**Chapter 11. HSM client dsmsys and
 dsmsopt option reference. 95**
 candidatesinterval 95
 checkfororphans. 96
 checkthresholds. 96
 compression 97
 defaultserver 98
 | errorlogname 98
 | errorprog 99
 | hsmdisableautomigdaemons 99
 | hsmdistributedrecall 100
 | hsmenableimmediatemigrate 101
 | hsmextobjdatr 101
 | hsmgroupedmigrate 102
 | hsmlogeventflags 103
 | hsmlogmax 103
 | hsmlogname 104
 | hsmlogretention 105
 | hsmlogsampleinterval 106
 | hsmmaxrecalltapedrives. 106
 | hsmmultiserver. 107
 | incl excl 108
 maxcandprocs 108
 maxmigrators 108
 maxrecalldaemons 109
 maxthresholdproc 110
 migfileexpiration 110
 migrateserver 111
 minmigfilesize 111
 minrecalldaemons 112
 optionformat 112
 reconcileinterval 113
 restoremigstate 113
 streamseq 115

**Chapter 12. HSM client command
 reference 117**
 Standard and short option formats 117
 Displaying command-line help 118
 Displaying file and file system information . . . 119
 HSM command summary 120
 dmkilld 121
 dsmmatr 122
 dsmautomig. 124
 dsmdf 126
 dsmdu 127
 dsmls 129
 dsmmigfs add and update 132
 | dsmmigfs addmultiserver, querymultiserver, and
 | removemultiserver 136
 | dsmmigfs deactivate, reactivate, and remove . 137
 | dsmmigfs globaldeactivate and globalreactivate 138
 | dsmmigfs help 139
 | dsmmigfs query 140
 | dsmmigfs rollback. 142

dsmmigfs sdrreset, enablefailover, and disablefailover	143	dsmMultiServerUpgrade.pl	167
dsmmigfs stop, start, and restart	144	dsmNextServername.pl	168
dsmmigfs takeover.	145	dsmreconcileGPFS.pl	168
dsmmighelp	146	dsmRemoveServer.pl	169
dsmmigquery	146	Appendix A. Control files in the	
dsmmigrate	149	.SpaceMan directory	171
dsmmigundelete	151	Space requirement for HSM control files	172
dsmmonitord	153	Appendix B. Accessibility features for	
dsmq	154	the Tivoli Storage Manager product	
dsmrecall	155	family.	175
dsmrecalld	158	Notices	177
dsmreconcile	158	Trademarks	179
dsmrm	161	Glossary	181
dsmrootd	161	Index	203
dsmscoutd	162		
dsmsetpw	163		
dsmwatchd	164		
 Chapter 13. HSM GPFS clients Perl			
 scripts	167		

Figures

1. start_HSM : sample script to start HSM	36		4. Ordered list file generation	82
2. stop_HSM : Stops HSM	37		5. Generation of list files	83
3. XML configuration file	40			

Tables

1.	Tivoli Storage Manager server publications	xii	17.	LAN-free data transfer options	34
2.	Tivoli Storage Manager storage agent publications	xii	18.	HSM agent options	41
3.	Tivoli Storage Manager client publications	xii	19.	dsmmigundelete command actions without the expiring option	71
4.	Tivoli Storage Manager data protection publications	xiii	20.	dsmmigundelete command actions with the expiring option	71
5.	IBM Tivoli Storage Manager troubleshooting and tuning publications	xiii	21.	HSM client determining recall modes	80
6.	AIX GPFS HSM client installation packages	7	22.	Automatic reconciliation tasks	85
7.	HSM AIX JFS2 installation packages	12	23.	Option format examples: options without values	118
8.	Linux GPFS available packages	14	24.	Option format examples: options with values	118
9.	Package names and descriptions	15	25.	HSM commands to display file and file system information	119
10.	Linux x86_64 GPFS uninstalling steps	17	26.	HSM command summary	120
11.	Space Management Options in Your Options Files	26	27.	Space management for locally and remotely mounted file systems	128
12.	HSM client options in the dsm.opt file	28	28.	Space management for locally and remotely mounted file systems	130
13.	HSM AIX attributes in a management class	29	29.	Control files stored in the .SpaceMan directory	171
14.	Assigning management classes to your files	30			
15.	Include and Exclude Statements	31			
16.	Environment Variables	33			

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

Publications for the IBM® Tivoli Storage Manager family of products are available online. The IBM Tivoli Storage Manager product family includes IBM Tivoli Storage FlashCopy® Manager, IBM Tivoli Storage Manager for Space Management, IBM Tivoli Storage Manager for Databases, and several other storage management products from IBM Tivoli.

To search all publications, go to the Tivoli Storage Manager information center at <http://publib.boulder.ibm.com/infocenter/tsminfo/v6r3>.

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/>.

Go to Tivoli Documentation Central to find information centers that contain official product documentation for current and previous versions of Tivoli products, including the Tivoli Storage Manager product family. You can find Tivoli Documentation Central at <https://www.ibm.com/developerworks/wikis/display/tivolidoccentral/Home>.

You can also order some related publications from the IBM Publications Center website. The website provides information about ordering publications from countries other than the United States. In the United States, you can order publications by calling 1-800-879-2755.

Tivoli Storage Manager publications

The following tables list the publications that make up the Tivoli Storage Manager library.

Table 1. Tivoli Storage Manager server publications

Publication title	Order number
<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
<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 Oracle Solaris Installation Guide</i>	GC23-9784
<i>IBM Tivoli Storage Manager for Oracle Solaris Administrator's Guide</i>	SC23-9772
<i>IBM Tivoli Storage Manager for Oracle 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 for z/OS Media Installation and User's Guide</i>	SC27-4018
<i>IBM Tivoli Storage Manager Upgrade and Migration Guide for V5 Servers</i>	GC27-4017
<i>IBM Tivoli Storage Manager Integration Guide for Tivoli Storage Manager FastBack[®]</i>	SC27-2828

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 Oracle 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

Table 3. Tivoli Storage Manager client publications (continued)

Publication title	Order number
<i>IBM Tivoli Storage Manager for Windows: Backup-Archive Clients Installation and User's Guide</i>	SC23-9792
<i>IBM Tivoli Storage Manager Using the Application Programming Interface</i>	SC23-9793
<i>IBM Tivoli Storage Manager for Space Management for UNIX and Linux: User's Guide</i>	SC23-9794
<i>IBM Tivoli Storage Manager HSM for Windows Administration Guide</i>	SC23-9795

Table 4. Tivoli Storage Manager data protection publications

Publication title	Order number
<i>IBM Tivoli Storage Manager for Databases: Data Protection for Microsoft SQL Server Installation and User's Guide</i>	GC27-4010
<i>IBM Tivoli Storage Manager for Databases: Data Protection for Oracle for UNIX and Linux Installation and User's Guide</i>	SC27-4019
<i>IBM Tivoli Storage Manager for Databases: Data Protection for Oracle for Windows Installation and User's Guide</i>	SC27-4020
<i>IBM Tivoli Storage Manager for Mail: Data Protection for Microsoft Exchange Server Installation and User's Guide</i>	GC27-4009
<i>IBM Tivoli Storage Manager for Mail: Data Protection for Lotus Domino® UNIX and Linux Installation and User's Guide</i>	SC27-4021
<i>IBM Tivoli Storage Manager for Mail: Data Protection for Lotus Domino for Windows Installation and User's Guide</i>	SC27-4022
<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 Virtual Environments Installation and User's Guide</i>	SC27-2898
<i>IBM Tivoli Storage Manager for Microsoft SharePoint Guide</i>	N/A

Table 5. IBM Tivoli Storage Manager troubleshooting and tuning publications

Publication title	Order number
<i>IBM Tivoli Storage Manager Problem Determination Guide</i>	GC23-9789
<i>IBM Tivoli Storage Manager Performance Tuning Guide</i>	GC23-9788
<i>IBM Tivoli Storage Manager Client Messages and Application Programming Interface Return Codes</i>	SC27-2878
<i>IBM Tivoli Storage Manager Server Messages and Error Codes</i>	SC27-2877
<i>IBM Tivoli Storage Manager for Mail: Data Protection for Microsoft Exchange Server Messages</i>	GC27-4011
<i>IBM Tivoli Storage Manager for Databases: Data Protection for Microsoft SQL Server Messages</i>	GC27-4012
<i>IBM Tivoli Storage Manager for Databases: Data Protection for Oracle Messages</i>	SC27-4014
<i>IBM Tivoli Storage Manager for Mail: Data Protection for Lotus Domino Messages</i>	SC27-4015

Table 5. IBM Tivoli Storage Manager troubleshooting and tuning publications (continued)

Publication title	Order number
IBM Tivoli Storage Manager for Enterprise Resource Planning: Data Protection for SAP Messages	SC27-4016

Note: You can find information about IBM System Storage® Archive Manager at http://publib.boulder.ibm.com/infocenter/tsminfo/v6r3/c_complydataretention_ovr.html.

Support information

You can find support information for IBM products from various sources.

Start at the IBM Support Portal: <http://www.ibm.com/support/entry/portal/>. You can select the products that you are interested in and search for a wide variety of relevant information.

Getting technical training

Information about Tivoli technical training courses is available online.

Go to the following websites to sign up for training, ask questions, and interact with others who use IBM storage products.

Tivoli software training and certification

Choose from instructor led, online classroom training, self-paced Web classes, Tivoli certification preparation, and other training options at <http://www.ibm.com/software/tivoli/education/>

Tivoli Support Technical Exchange

Technical experts share their knowledge and answer your questions in webcasts at http://www.ibm.com/software/sysmgmt/products/support/supp_tech_exch.html.

Storage Management community

Interact with others who use IBM storage management products at <http://www.ibm.com/developerworks/servicemanagement/sm/index.html>

Global Tivoli User Community

Share information and learn from other Tivoli users throughout the world at <http://www.tivoli-ug.org/>.

IBM Education Assistant

View short "how to" recordings designed to help you use IBM software products more effectively at <http://publib.boulder.ibm.com/infocenter/ieduasst/tivv1r0/index.jsp>

Searching knowledge bases

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

Begin by searching the Tivoli Storage Manager Information Center at <http://publib.boulder.ibm.com/infocenter/tsminfo/v6r3>. From this website, you can search the current Tivoli Storage Manager documentation.

Searching the Internet

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

To search multiple Internet resources, go to the IBM support website at <http://www.ibm.com/support/entry/portal/>.

You can search for information without signing in. Sign in using your IBM ID and password if you want to customize the site based on your product usage and information needs. If you do not already have an IBM ID and password, click **Sign in** at the top of the page and follow the instructions to register.

From the support website, you can search various resources including:

- IBM technotes
- IBM downloads
- IBM Redbooks® publications
- IBM Authorized Program Analysis Reports (APARs)

Select the product and click **Downloads** to search the APAR list.

If you still cannot find a solution to the problem, you can search forums and newsgroups on the Internet for the latest information that might help you find problem resolution.

An independent user discussion list, ADSM-L, is hosted by Marist College. You can subscribe by sending an email to listserv@vm.marist.edu. The body of the message must contain the following text: `SUBSCRIBE ADSM-L your_first_name your_family_name`.

To share your experiences and learn from others in the Tivoli Storage Manager and Tivoli Storage FlashCopy Manager user communities, go to the following wikis:

Tivoli Storage Manager wiki

<http://www.ibm.com/developerworks/wikis/display/tivolistoragemanager>

Tivoli Storage FlashCopy Manager wiki

[https://www.ibm.com/developerworks/mydeveloperworks/wikis/home/wiki/Tivoli Storage FlashCopy Manager](https://www.ibm.com/developerworks/mydeveloperworks/wikis/home/wiki/Tivoli%20Storage%20FlashCopy%20Manager)

Using IBM Support Assistant

IBM Support Assistant is a complimentary software product that can help you with problem determination. It is available for some Tivoli Storage Manager and Tivoli Storage FlashCopy Manager products.

To learn about which products are supported, go to the IBM Support Assistant download web page at <http://www.ibm.com/software/support/isa/download.html>.

IBM Support Assistant helps you gather support information when you must 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

You can find more information at the IBM Support Assistant website:

<http://www.ibm.com/software/support/isa/>

You can also install the stand-alone IBM Support Assistant application on any workstation. You can then enhance the application by installing product-specific plug-in modules for the IBM products that you use. Find add-ons for specific products at <http://www.ibm.com/support/docview.wss?uid=swg27012689>.

Finding product fixes

A product fix to resolve your problem might be available from the IBM software support website.

You can determine what fixes are available by checking the IBM software support website at <http://www.ibm.com/support/entry/portal/>.

- If you previously customized the site based on your product usage:
 1. Click the link for your product, or a component for which you want to find a fix.
 2. Click **Downloads**, and then click **Fixes by version**.
- If you have not customized the site based on your product usage, click **Downloads** and search for your product.

Receiving notification of product fixes

You can receive notifications about fixes, flashes, upgrades, and other news about IBM products.

To sign up to receive notifications about IBM products, follow these steps:

1. From the support page at <http://www.ibm.com/support/entry/portal/>, click **Sign in to create, manage, or view your subscriptions** in the **Notifications** pane.
2. Sign in using your IBM ID and password. If you do not have an ID and password, click **register now** and complete the registration process.
3. Click **Manage all my subscriptions** in the **Notifications** pane.
4. Click the **Subscribe** tab and then click **Tivoli**.
5. Select the products for which you want to receive notifications and click **Continue**.

6. Specify your notification preferences and click **Submit**.

Contacting IBM Software Support

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

To obtain help from IBM Software Support, complete the following steps:

1. Ensure that you have completed the following prerequisites:
 - a. Set up a subscription and support contract.
 - b. Determine the business impact of your problem.
 - c. Describe your problem and gather background information.
2. Follow the instructions in “Submitting the problem to IBM Software Support” on page xviii.

Setting up a subscription and support contract

Set up a subscription and support 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, IBM Tivoli, Lotus®, and Rational® products, as well as IBM DB2® and IBM WebSphere® products that run on Microsoft Windows or on operating systems such as AIX or Linux), enroll in IBM Passport Advantage® in one of the following ways:

- **Online:** Go to the Passport Advantage website at <http://www.ibm.com/software/lotus/passportadvantage/>, click **How to enroll**, and follow the instructions.
- **By telephone:** You can call 1-800-IBMSERV (1-800-426-7378) in the United States. For the telephone 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**.

Determining the business impact

When you report a problem to IBM, you are asked to supply a severity level. Therefore, you must 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.

Describing the problem and gathering 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 re-created? 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 using a workaround for this problem? If so, be prepared to explain it when you report the problem.

Submitting the problem to IBM Software Support

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

Online

Go to the IBM Software Support website at [http://www.ibm.com/support/entry/portal/Open_service_request/Software/Software_support_\(general\)](http://www.ibm.com/support/entry/portal/Open_service_request/Software/Software_support_(general)). Sign in to access IBM Service Requests and enter your information into the problem submission tool.

By telephone

For the telephone number to call in your country, go to the IBM Software Support Handbook at <http://www14.software.ibm.com/webapp/set2/sas/f/handbook/home.html> and click **Contacts**.

Conventions used in this publication

This publication uses the following typographical conventions:

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

Tivoli Storage Manager for Space Management version 6.3 updates

Tivoli Storage Manager for Space Management is updated for version 6.3.

The following features are included in IBM Tivoli Storage Manager for Space Management Version 6.3:

- The Hierarchical Storage Manager (HSM) client manages space on the following file systems:
 - General Parallel File System (GPFS™) on AIX® cluster
 - General Parallel File System (GPFS) on Linux x86_64 cluster
 - Enhanced Journaled File System (JFS2) on an AIX workstation
- The HSM logging facility. The HSM logging facility provides records of file migration and recall. These records are composed of date, time, Data Management Application Programming Interface (DMAPI) handle, and result (success or failure). The log also records file system events such as threshold migration, reconfiguration, and reconciliation. This log does not contain error information.
- The name of the error log can be specified for each command issued by passing the **errorlogname** option when issuing the command.
- GPFS integration for threshold migration is enhanced with the **hsmdisableautomigdaemons** option.
- Multiple Tivoli Storage Manager servers can manage a single file system. This feature is supported with AIX GPFS and Linux GPFS HSM clients.
- The **dsmrecall** command is enhanced with the option to order a file list so that tape access is optimized during a recall.
- The **maxreconcileproc** option is no longer valid.
- The **optionformat** option is a deprecated option. It is supported in version 6.3, but will not be supported in future versions. In future versions, all options must be specified with the standard format.

Related concepts

Chapter 1, “HSM client for UNIX and Linux overview,” on page 1

“Optimized tape recall processing” on page 80

“Managing a file system with multiple Tivoli Storage Manager servers” on page 53

Related tasks

“Editing the options file dsm.sys” on page 25

Related reference

“Setting HSM logging” on page 41

Chapter 1. HSM 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 space-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 a root user, can do the following:

- 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 process 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. The IBM Tivoli Storage Manager backup-archive client is used 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. The backup-archive client is used 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.

Related concepts

Chapter 2, “Installing the HSM for UNIX and Linux client,” on page 5

Chapter 5, “Migrating files,” on page 61

Chapter 6, “Backing up and restoring file systems,” on page 67

Chapter 7, “Recalling migrated files,” on page 77

Chapter 8, “Reconciling file systems,” on page 85

Related reference

“Tivoli Storage Manager for Space Management version 6.3 updates” on page xix

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, 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_64 GPFS, the process receives an out-of-space (ENOSPC) return code and stops.

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

For information about configuring GPFS integration with the HSM client, see the Tivoli field guide *TSM for Space Management for UNIX-GPFS Integration Part I: Policy-driven Threshold Migration* at <https://www.ibm.com/support/docview.wss?uid=swg27018848>.

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.

Related concepts

Chapter 5, “Migrating files,” on page 61

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.

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.

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

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

Partial file recall mode recalls a portion of a migrated file and is valid for AIX GPFS and Linux x86_64 GPFS only. A partial recall avoids recalling an entire file, when only a small portion of the file is required by an application. The HSM client intercepts a read request for a file configured for partial file recall. The HSM client then calculates which portion of the file to recall based on the offsets contained in the read request. This results in time and disk space savings, since only a portion of the file is recalled.

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

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

Related concepts

Chapter 7, “Recalling migrated files,” on page 77

“Selective recall processing” on page 77

“Optimized tape recall processing” on page 80

“Partial file recall mode” on page 78

“Streaming recall mode” on page 79

Reconciling file systems overview

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

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

Related concepts

Chapter 8, “Reconciling file systems,” on page 85

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.

The HSM client manages space on the following file systems:

- General Parallel File System (GPFS) on AIX cluster
- General Parallel File System (GPFS) on Linux x86_64 cluster
- Enhanced Journaled File System (JFS2) on an AIX workstation

Related concepts

Chapter 1, “HSM client for UNIX and Linux overview,” on page 1

“Assigning management classes to files” on page 29

“Registering your workstation with a Tivoli Storage Manager server” on page 18

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 and set up the Tivoli Storage Manager 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.

For more information about the Tivoli Storage Manager backup-archive clients, see the *IBM Tivoli Storage Manager for UNIX and Linux Backup-Archive Clients Installation and User's Guide*.

Note: See the *IBM Tivoli Storage Manager for UNIX and Linux Backup-Archive Clients Installation and User's Guide* for supported languages and locales.

Pre-existing files in the installation directory can be deleted

The `/usr/tivoli/tsm/client/hsm/bin` directory is one of the base directories into which the 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. Before installing anything, read all of these general requirements. Read your system installation overview, specific environment requirements, and each installation step to ensure that you are ready to install the HSM client.
2. Follow your system-specific installation procedure.
3. When you complete the installation, register your workstation as a node with a Tivoli Storage Manager server.
4. Modify the `dsm.sys` and `dsm.opt` configuration files. Both files are located in the following directory:
 - For AIX: `/usr/tivoli/tsm/client/ba/bin`
 - For Linux: `/opt/tivoli/tsm/client/ba/bin`

Reinstallation or upgrade overview

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

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

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:
 - The Tivoli Storage Manager API, Tivoli Storage Manager backup-archive client, and Tivoli Storage Manager HSM client
 - The Oracle Java™ Runtime Environment (JRE) version 6.
- 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 file. In order to verify that you have the right JRE installed, issue the `java -version` command from a UNIX terminal window.
- If you do not have the requisite version of the JRE, you must download the required version of the JRE.

HSM client limitations and considerations

- The HSM AIX GPFS client is not compatible with the backup-archive client for JFS2. Journal-based backup is not supported when the HSM AIX GPFS client is installed.
- The HSM GPFS clients manages only file systems belonging to the local (home) GPFS cluster; it does not manage remotely mounted file systems.

- In a GPFS environment, a small file that is less than the size of a GPFS block size can become larger after an HSM migration. GPFS adds meta information to the file during the migration. Because another block on the file system is allocated for the meta information, the space allocated for the file increases. If a file system is filled to its maximum capacity with many small files, it can run out of space during the file migration.
- The HSM client is not supported on JFS2 Encrypted File Systems (EFS).
- Extended attributes (EA_{v2}) are supported by the HSM client. However, if you delete a stub file and use the **dsmmigundelete** command to recreate the stub file, the extended attributes are not restored.

Related concepts

Chapter 3, “Configuring the HSM client,” on page 23

“Space requirement for HSM control files” on page 172

AIX GPFS HSM client installation overview

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

There are several installation limitations for the HSM AIX GPFS client:

- 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 must remove the JFS2 clients.
- On AIX 6.1 and 7.1, the HSM client can be installed in the global partition and supports transparent recall for both global and local workstation partitions (WPARs). Using HSM commands from a local WPAR is not supported. You cannot install the HSM client in a local WPAR.

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

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

Table 6 indicates the packages available on the installation media in the `/usr/sys/inst.images` directory:

Table 6. AIX GPFS HSM client installation packages

Package	Installs	Into this directory
<code>tivoli.tsm.client.ba64.gpfs</code>	The backup-archive client for AIX GPFS	<code>/usr/tivoli/tsm/client/ba/bin</code>
<code>tivoli.tsm.client.hsm.gpfs</code>	The HSM client for AIX GPFS	<code>/usr/tivoli/tsm/client/hsm/bin</code>

Table 6. AIX GPFS HSM client installation packages (continued)

Package	Installs	Into this directory
tivoli.tsm.client.api.64bit	The API for AIX	/usr/tivoli/tsm/client/api/bin

For an initial installation, follow these steps:

1. If you want the GPFS policy engine to control automatic migration, you can disable the automatic migration daemons **dsmonitord** and **dsmscoutd**. Disabling these daemons conserves system resources. To disable the automatic migration daemons, start this command in a shell: `export hsminstallmode=scoutfree`.
For information about configuring GPFS integration with the HSM client, see the Tivoli field guide *TSM for Space Management for UNIX-GPFS Integration Part I: Policy-driven Threshold Migration* at <https://www.ibm.com/support/docview.wss?uid=swg27018848>.
2. Edit the `dsm.opt` and `dsm.sys` files installed with the backup-archive client to configure the HSM client.
3. Install the HSM client on each HSM node. For AIX clients, see “Installing the AIX HSM client” on page 9. For Linux clients, see “Installing the Linux x86_64 GPFS client” on page 15.
4. Make sure that after installation, the **dsmreca1ld** daemon is running on at least one node.
5. Enable the Data Management Application Programming Interface (DMAPI) for GPFS for all file systems to which you plan to add space management. Enable DMAPI only once for each file system.
 - a. Unmount all GPFS file systems on all nodes within the GPFS cluster to which you plan to add space management.
 - b. Activate DMAPI management for the GPFS file systems with the following command:
AIX: `mmchfs device -z yes`
LINUX: `mmchfs -z yes`
See the *GPFS: Administration and Programming Reference* for information about GPFS commands.
 - c. Remount all GPFS file systems on all nodes within the GPFS cluster.
The HSM daemons detect the initial state of each node and assign all nodes an instance number in relation to the GPFS cluster definition.
6. On the HSM owner nodes, add space management to each GPFS file system with the **dsmmigfs** command.
7. Use the **dsmmigfs enablefailover** command. to enable failover of space management on the owner and source cluster nodes that participate in distributed space management.

Related tasks

“Editing the options files using the backup-archive client” on page 24

Related reference

“Adding space management to WPARS on AIX 6.1 and 7.1” on page 46

“**dsmmigfs add** and **update**” on page 132

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

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 General Parallel File System (GPFS) HSM client requires the RS/6000® 64-bit pSeries® or compatible hardware.

Software requirements

The AIX GPFS HSM client requires the following software:

- AIX 5L™ 6.1 PPC in 64-bit kernel mode or AIX 5L 7.1 PPC in 64-bit kernel mode
- GPFS 3.4

Installing the AIX HSM client

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

Follow these steps from the AIX command line or DVD. These steps install the AIX HSM General Parallel File System (GPFS) client or the AIX HSM Enhanced Journaled File System (JFS2) client.

1. Optional: Specify **hsminstallmode=scoutfree**.

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

For information about configuring GPFS integration with the HSM client, see the Tivoli field guide *TSM for Space Management for UNIX-GPFS Integration Part I: Policy-driven Threshold Migration* at <https://www.ibm.com/support/docview.wss?uid=swg27018848>.

2. Issue the **smitty install** command. If you are installing from a hard disk directory, remove the `.toc` file from that directory before issuing the **smitty** command.
3. Select **Install and Update Software** and press Enter.
4. Select **Install and Update from ALL Available Software** and press Enter.
5. In the **INPUT device / directory for software** field, select one of the following options and press Enter.
 - a. `/usr/sys/inst.images` (directory of installation images)
 - b. `/dev/cd0` (installation from DVD)
6. In the **SOFTWARE to Install** field press F4.
7. 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 the following:
 - Tivoli Storage Manager application programming interface
 - Tivoli Storage Manager backup-archive client common files
 - Tivoli Storage Manager backup-archive client base files
8. Select the options that you want and press Enter to begin the installation.
9. If needed, unmount the DVD drive.

After installing, you can register your workstation with a Tivoli Storage Manager server.

Related concepts

“Registering your workstation with a Tivoli Storage Manager server” on page 18

Upgrading the AIX GPFS HSM client

You can upgrade the HSM GPFS client, and then upgrade from General Parallel File System (GPFS) Version 2.2 or lower.

Before upgrading the HSM client, read the HSM GPFS client installation overview.

If you are upgrading GPFS from version 2.2 or lower, follow the GPFS migration procedure from *GPFS: Concepts, Planning, and Installation Guide*.

If you are upgrading GPFS to version 3.4 or later, upgrade all cluster nodes to the same version and set the option **dmapihandle size=32**. If these tasks are not done, the HSM client might fail with the following errors:

- The **dsmautomic** 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.file_system` directory is created with a size of 0.

To check if the **dmapihandle size** parameter is set to the correct value, issue the following command: `mmfsadm dump config | grep -i dmapiHandleSize`

If the output from the command is `dmapiHandleSize 16`, set option **dmapihandle size=32**.

To upgrade the HSM GPFS client, perform the following steps:

1. Ensure that all HSM nodes (owner and backup) are in a consistent state and that all space-managed file systems are mounted on all HSM nodes.
2. Globally deactivate HSM on every node by issuing the command: **dsmmigfs globaldeactivate .**
3. Disable failover on every node by issuing the command: **dsmmigfs disablefailover .**
4. To figure out which node owns which cluster file system, issue the **dsmmigfs q -d** command on one of the nodes.
5. Remove Tivoli Storage Manager for Space Management from all nodes. You must not upgrade with a command like **rpm -U**. You must uninstall the old version and then install the new version.
6. Install Tivoli Storage Manager for Space Management on all nodes.
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 also run on these nodes, which are designated for failover feature.

Related concepts

“AIX GPFS HSM client installation overview” on page 7

Related reference

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

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

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

“**dsmmigfs query**” on page 140

Uninstalling the AIX GPFS HSM client

Follow this procedure to uninstall the AIX GPFS HSM client.

Before uninstalling the HSM client, read the AIX GPFS HSM client installation overview (“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 N0 and press Enter.

Related concepts

“AIX GPFS HSM client installation overview” on page 7

Related reference

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

AIX JFS2 HSM client installation overview

Before installing the HSM client on AIX Enhanced Journaled File System (JFS2), you must review the general requirements, system-specific requirements, and review all installation steps.

There are several installation limitations for the HSM AIX JFS2 client:

- The HSM AIX JFS2 client is not compatible with the HSM AIX General Parallel File System (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 must remove the GPFS clients.
- On AIX 6.1 and 7.1, the HSM client can be installed in the global partition and supports transparent recall for both global and local workstation partitions (WPARs). Using HSM commands from a local WPAR is not supported. You cannot install the HSM client in a local WPAR.

Table 7 on page 12 lists the packages that are available on the installation media:

Table 7. HSM AIX JFS2 installation packages

Package	Installs	Into the following directory
tivoli.tsm.client.ba.64bit	The backup-archive client for AIX	/usr/tivoli/tsm/client/ba/bin
tivoli.tsm.client.hsm.jfs2	The HSM client for AIX JFS2 64-bit	/usr/tivoli/tsm/client/hsm/bin
tivoli.tsm.client.api.64bit	The API for AIX	/usr/tivoli/tsm/client/api/bin
tivoli.tsm.client.msg.lang	The national language messages for the HSM client. American English messages are already included in the HSM client.	/usr/tivoli/tsm/client/lang/bin

Related reference

“**dsmmigfs add** and **update**” on page 132

“Adding space management to WPARS on AIX 6.1 and 7.1” on page 46

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

Software requirements

The HSM AIX JFS2 client requires the following software:

- AIX 5L 6.1 in 64-bit kernel mode or AIX 5L 7.1 in 64-bit kernel mode

Installing the AIX HSM client

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

Follow these steps from the AIX command line or DVD. These steps install the AIX HSM General Parallel File System (GPFS) client or the AIX HSM Enhanced Journaled File System (JFS2) client.

1. Optional: Specify **hsminstallmode=scoutfree**.

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

For information about configuring GPFS integration with the HSM client, see the Tivoli field guide *TSM for Space Management for UNIX-GPFS Integration Part I: Policy-driven Threshold Migration* at <https://www.ibm.com/support/docview.wss?uid=swg27018848>.

2. Issue the **smitty install** command. If you are installing from a hard disk directory, remove the `.toc` file from that directory before issuing the **smitty** command.
3. Select **Install and Update Software** and press Enter.
4. Select **Install and Update from ALL Available Software** and press Enter.
5. In the **INPUT device / directory for software** field, select one of the following options and press Enter.

- a. /usr/sys/inst.images (directory of installation images)
 - b. /dev/cd0 (installation from DVD)
6. In the **SOFTWARE to Install** field press F4.
 7. 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 the following:
 - Tivoli Storage Manager application programming interface
 - Tivoli Storage Manager backup-archive client common files
 - Tivoli Storage Manager backup-archive client base files
 8. Select the options that you want and press Enter to begin the installation.
 9. If needed, unmount the DVD drive.

After installing, you can register your workstation with a Tivoli Storage Manager server.

Related concepts

“Registering your workstation with a Tivoli Storage Manager server” on page 18

Linux x86_64 GPFS HSM client installation overview

Before installing the HSM client on Linux x86_64 General Parallel File System (GPFS), review both the general and the system-specific requirements. If you are installing the product for the first time, use the steps for an initial installation. Otherwise, use the steps for an upgrade.

Note:

- HSM cluster installations are certified on IBM Linux Cluster 1350. See the IBM Redbooks: *Linux Clustering with CSM and GPFS*.
- Also see the recommendations provided with the IBM GPFS for Linux on x86_64 architecture.

When you install the HSM client on Linux x86_64 GPFS file systems, the installation process does the following:

- Stops any space management daemons that are running.
- Removes any statement from the /etc/inittab file that starts the **dsmwatchd** daemon at system startup.
- Removes any statement from the /var/mmfs/etc/gpfsready script file that loads the other space management daemons at GPFS system startup.
- Extracts the HSM modules.
- Adds a statement to the /etc/inittab file that starts the **dsmwatchd** daemon at system startup.
- Adds a statement to the /var/mmfs/etc/gpfsready script file that starts the other space management daemons at GPFS system startup.
- Starts the space management daemons.
- 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 8 on page 14 lists the packages available on the installation media:

Table 8. Linux GPFS available packages

Package	Installs	Into this directory
TIVsm-API64.x86_64.rpm	The API for Linux x86_64 (64-bit only)	/opt/tivoli/tsm/client/api/bin64
TIVsm-BA.x86_64.rpm	The Tivoli Storage Manager backup-archive client (command line), the administrative client (command line), and the Web backup-archive client (64-bit only)	/opt/tivoli/tsm/client/ba/bin
TIVsm-HSM.x86_64.rpm	The HSM client for Linux x86_64 (64-bit only)	/opt/tivoli/tsm/client/hsm/bin
TIVsm-msg.lang.x86_64.rpm	The message package for language (where lang is the language code)	/opt/tivoli/tsm/client/lang

For an initial installation, follow these steps:

1. If you want the GPFS policy engine to control automatic migration, you can disable the automatic migration daemons **dsmonitord** and **dsmscoutd**. Disabling these daemons conserves system resources. To disable the automatic migration daemons, start this command in a shell: `export hsminstallmode=scoutfree`.
For information about configuring GPFS integration with the HSM client, see the Tivoli field guide *TSM for Space Management for UNIX-GPFS Integration Part I: Policy-driven Threshold Migration* at <https://www.ibm.com/support/docview.wss?uid=swg27018848>.
2. Edit the `dsm.opt` and `dsm.sys` files installed with the backup-archive client to configure the HSM client.
3. Install the HSM client on each HSM node. For AIX clients, see “Installing the AIX HSM client” on page 9. For Linux clients, see “Installing the Linux x86_64 GPFS client” on page 15.
4. Make sure that after installation, the **dsmsrecalld** daemon is running on at least one node.
5. Enable the Data Management Application Programming Interface (DMAPI) for GPFS for all file systems to which you plan to add space management. Enable DMAPI only once for each file system.
 - a. Unmount all GPFS file systems on all nodes within the GPFS cluster to which you plan to add space management.
 - b. Activate DMAPI management for the GPFS file systems with the following command:
AIX: `mmchfs device -z yes`
LINUX: `mmchfs -z yes`
See the *GPFS: Administration and Programming Reference* for information about GPFS commands.
 - c. Remount all GPFS file systems on all nodes within the GPFS cluster.
The HSM daemons detect the initial state of each node and assign all nodes an instance number in relation to the GPFS cluster definition.
6. On the HSM owner nodes, add space management to each GPFS file system with the **dsmmigfs** command.
7. Use the **dsmmigfs enablefailover** command to enable failover of space management on the owner and source cluster nodes that participate in distributed space management.

Related tasks

“Editing the options files using the backup-archive client” on page 24

Related reference

“**dsmmigfs add** and **update**” on page 132

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

Linux x86_64 GPFS HSM client environment requirements

Before you install the HSM client on a Linux x86_64 General Parallel File System (GPFS), your workstation must meet minimum communication, hardware, software requirements.

Software requirements

The HSM client requires the following software in order to run with the following distributions:

- GPFS 3.4 or GPFS 3.5 on 64-bit SUSE Linux Enterprise Server 10 or 11 or Red Hat Enterprise Linux 5.0 or 6.0

Note: X Window 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_64 GPFS client

You can install the HSM client on Linux x86_64 General Parallel File System (GPFS) using the product DVD.

Follow these steps:

1. Log in as the root user and mount the DVD to the `/cdrom` directory if applicable.
2. Change directory to the installation packages directory. On the DVD:
`cdrom/tsmcli/linux86`
3. Optional: Specify `hsminstallmode=scoutfree`.

If you want the GPFS policy engine to control automatic migration, you can disable the automatic migration daemons `dsmmonitor` and `dsmscoutd`. Disabling these daemons conserves system resources. To disable the automatic migration daemons, start this command in a shell: `export hsminstallmode=scoutfree`.

For information about configuring GPFS integration with the HSM client, see the Tivoli field guide *TSM for Space Management for UNIX-GPFS Integration Part I: Policy-driven Threshold Migration* at <https://www.ibm.com/support/docview.wss?uid=swg27018848>.

4. Install the Tivoli Storage Manager clients in the order that is presented in Table 9. During installation, these packages are installed in unique directories. Press `y` to each question when prompted.
5. Restart your system when the installation is complete.

Table 9. Package names and descriptions

To install	Issue this command
API	Install the API files in the <code>/opt/tivoli/tsm/client/api/bin64</code> directory. <code>rpm -i TIVsm-API64.x86_64.rpm</code>

Table 9. Package names and descriptions (continued)

To install	Issue this command
Backup-archive client (CLI and GUI) web client Administrative client	<p>Install the backup-archive client (command line, GUI, and web client) in the <code>/opt/tivoli/tsm/client/ba/bin</code> directory.</p> <pre>rpm -i TIVsm-BA.x86_64.rpm</pre> <p>To circumvent the dependence check, you can use the <code>--nodeps</code> option, but then you must check the dependencies manually.</p> <p>To install the backup-archive client (command line, GUI, and web client) and administrative client in the <code>/opt/tivoli/tsm/client/ba/bin</code> directory:</p> <pre>rpm -i TIVsm-BA.x86_64.rpm</pre>
HSM client	<p>To install the following HSM client files in the <code>/opt/tivoli/tsm/client/hsm/bin</code> directory.</p> <ul style="list-style-type: none"> • HSM client commands • Space management agent (hsmagent) • IBM Tivoli Enterprise Space Management Console (HSM GUI) <pre>rpm -i TIVsm-HSM.x86_64.rpm</pre>

Related reference

“`hsmdisableautomigdaemons`” on page 99

Upgrading the Linux x86_64 GPFS HSM client

You can upgrade the HSM GPFS client, and then upgrade from General Parallel File System (GPFS) Version 2.2 or lower.

Before upgrading the HSM client, read the HSM GPFS client installation overview.

If you are upgrading GPFS from version 2.2 or lower, follow the GPFS migration procedure from *GPFS: Concepts, Planning, and Installation Guide*.

If you are upgrading GPFS to version 3.4 or later, upgrade all cluster nodes to the same version and set the option `dmapihandle=32`. If these tasks are not done, the HSM client might fail with the following errors:

- The `dsmautomic` 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.file_system` directory is created with a size of 0.

To check if the `dmapihandle` parameter is set to the correct value, issue the following command: `mmfsadm dump config | grep -i dmapihandle`

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

To upgrade the HSM GPFS client, perform the following steps:

1. Ensure that all HSM nodes (owner and backup) are in a consistent state and that all space-managed file systems are mounted on all HSM nodes.
2. Globally deactivate HSM on every node by issuing the command: `dsmmigfs globaldeactivate`.

3. Disable failover on every node by issuing the command: **dsmmigfs disablefailover** .
4. To figure out which node owns which cluster file system, issue the **dsmmigfs q -d** command on one of the nodes.
5. Remove Tivoli Storage Manager for Space Management from all nodes. You must not upgrade with a command like **rpm -U**. You must uninstall the old version and then install the new version.
6. Install Tivoli Storage Manager for Space Management on all nodes.
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 also run on these nodes, which are designated for failover feature.

Related concepts

“Linux x86_64 GPFS HSM client installation overview” on page 13

Related reference

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

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

“**dsmmigfs query**” on page 140

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

Uninstalling the Linux x86_64 GPFS HSM client

You must follow a specific procedure to uninstall the Linux 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 space-managed file systems should be avoided during software installation. Access to migrated files during the upgrade is not possible, because the system would be suspended while trying to read a file stub.

Use Table 10 to know which steps to run to uninstall the HSM client.

Table 10. Linux x86_64 GPFS uninstalling steps

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

After uninstalling, you can install the latest version.

Related concepts

“Linux x86_64 GPFS HSM client installation overview” on page 13

Related reference

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

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
```

Related reference

“`dsmsetpw`” on page 163

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 so you can remotely administer multiple HSM clients from a single point. Download the Windows package from <ftp://public.dhe.ibm.com/storage/tivoli-storage-management/>.

Before using HSM GUI, check the following:

1. Java Runtime Environment (JRE) Version 6 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.
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.
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 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 manage 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. Configuring the HSM client

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

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

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

During installation of the 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 Linux x86\_64: `/opt/tivoli/tsm/client/ba/bin`

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

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

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 UNIX and Linux backup-archive clients:

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

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

For information about configuring GPFS integration with the HSM client, see the Tivoli field guide *TSM for Space Management for UNIX-GPFS Integration Part I: Policy-driven Threshold Migration* at <https://www.ibm.com/support/docview.wss?uid=swg27018848>.

#### **Related concepts**

Chapter 11, “HSM client `dsm.sys` and `dsm.opt` option reference,” on page 95

---

## **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 F1, 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.

#### **Related reference**

“`dsmmigquery`” on page 146

---

## **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 client 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.

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

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

## Editing the options file `dsm.sys`

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

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

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

Options are processed following this order:

1. Options that are defined on the server with server-enforced client options (the client cannot override the value)
2. Options that are entered locally on the command line
3. Options that are defined on the server for a schedule 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 26 provides a brief description of each space management option that you can set in your `dsm.sys` file.

### Note:

- Except for **errorprog**, place all space management options at the beginning of the `dsm.sys` file before any server stanzas that you define.
- You can specify both a default server and a migration server in the `dsm.sys` file. Use the **defaultserver** option and the **migrateserver** option. If you do not specify migration server and default server, the server that you specify in the first stanza of the `dsm.sys` file becomes the default server.
- If you specify a migration server with the **migrateserver** option, it overrides the server that you specify with the **defaultserver** option.

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

```
DEFAULTServer server1
MIGRATEServer server2
CHECKThresholds 2
CANDIDATESInterval 12
MAXCANDprocs 5
RECOncileinterval 1
MAXRECOncileproc 5
MAXThresholdproc 5
MINMIGFILESize 8192
MIGFILEEXPIration 10
MINRECALLdaemons 5
MAXRECALLdaemons 15
CHECKFororphans no
MAXMIGRators 1
KERNelmessages no
OVERLAPRECALL no
Servername server1
 COMMmethod TCPip
 TCPPort 1500
 TCPServeraddress server3.almaden.ibm.com
 Passwordaccess generate
 Mailprog /usr/bin/xsend root
 Groups system tsm
```

|                   |                           |
|-------------------|---------------------------|
| Users             | steiner chron wang nguyen |
| Incl excl         | /adm/tsm/backup.excl      |
| ERRORProg         | /bin/cat                  |
| <b>Servername</b> | <b>server2</b>            |
| COMMethod         | SNA1u6.2                  |
| PARTner1uname     | raptor                    |
| TPname            | apcdcl                    |
| CPICM0dename      | appc                      |
| Passwordaccess    | generate                  |
| Mailprog          | /usr/bin/xsend root       |
| Groups            | system tsm                |
| Users             | sullivan tang stewart     |
| Incl excl         | /adm/tsm/migrate.excl     |
| ERRORProg         | /bin/cat                  |

Table 11. Space Management Options in Your Options Files

| Option                      | Default | Description                                                                                                                                                                                                                                                                                               |
|-----------------------------|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>candidatesinterval</b>   | 1       | Specifies how often the <b>dsmscoutd</b> 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 “ <b>candidatesinterval</b> ” on page 95.                                                                 |
| <b>checkfororphans</b>      | NO      | Specifies whether the <b>ds mreconcile</b> command checks for orphans. Specify YES or NO. See “ <b>checkfororphans</b> ” on page 96.                                                                                                                                                                      |
| <b>checkthresholds</b>      | 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 “ <b>checkthresholds</b> ” on page 96.                                                                                            |
| <b>compression</b>          | NO      | Compresses files before you send them to the Tivoli Storage Manager server. Specify YES or NO. See “ <b>compression</b> ” on page 97.                                                                                                                                                                     |
| <b>defaultserver</b>        |         | Specifies the name of the default Tivoli Storage Manager server. This server is contacted for space management services if you do not specify the <b>migrateserver</b> option. See “ <b>defaultserver</b> ” on page 98.                                                                                   |
| <b>errorprog</b>            |         | Specifies a program to which you want to send a message if a severe error occurs during space management processing. Place this option after all space management options for each server that you define in your <b>dsm.sys</b> file. See “ <b>errorprog</b> ” on page 99.                               |
| <b>hsmextobjidattr</b>      | NO      | Specifies whether the HSM client assigns an object ID as a Data Management Application Programming Interface (DMAPI) extended attribute (extObjId) to premigrated or migrated files. The DMAPI attribute is required for two-way orphan check reconciliation. See “ <b>hsmextobjidattr</b> ” on page 101. |
| <b>hsmlogeventflags</b>     |         | Specifies which events are added to the HSM log. You can specify file events, file system events, or system events. See “ <b>hsmlogeventflags</b> ” on page 103.                                                                                                                                          |
| <b>hsmlogmax</b>            | 0       | Specifies the maximum age for retaining log entries in the HSM log file. Older entries are pruned. See “ <b>hsmlogmax</b> ” on page 103.                                                                                                                                                                  |
| <b>hsmlogname</b>           |         | Specifies the name of the HSM log file. See “ <b>hsmlogname</b> ” on page 104.                                                                                                                                                                                                                            |
| <b>hsmlogretention</b>      | N,D     | Specifies how many days to maintain HSM log entries before pruning, and whether to save the pruned entries. See “ <b>hsmlogretention</b> ” on page 105.                                                                                                                                                   |
| <b>hsmlogsampleinterval</b> | 3600    | Specifies how often to gather file system data for the HSM log file. See “ <b>hsmlogsampleinterval</b> ” on page 106.                                                                                                                                                                                     |
| <b>incl excl</b>            |         | Specifies the file name and path of your include-exclude file. See “Modifying the include-exclude options file” on page 31. See “ <b>incl excl</b> ” on page 108.                                                                                                                                         |

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

| Option                   | Default  | Description                                                                                                                                                                                                                                                                                                                                |
|--------------------------|----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>maxcandprocs</b>      | 5        | Specifies the number of parallel processes in the scout daemons which can scan file systems. The range of values is 1 through 20. See “ <b>maxcandprocs</b> ” on page 108.                                                                                                                                                                 |
| <b>maxmigrators</b>      | 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 must be made in the Tivoli Storage Manager server configuration to update the HSM node <b>MAXNUMMP</b> value. See “ <b>maxmigrators</b> ” on page 108. |
| <b>maxrecalldaemons</b>  | 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 <b>maxrecalldaemons</b> value is less than <b>minrecalldaemons</b> value, the default is set for both options. See “ <b>maxrecalldaemons</b> ” on page 109.                   |
| <b>maxthresholdproc</b>  | 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 “ <b>maxthresholdproc</b> ” on page 110.                                                                                                                                                  |
| <b>migfileexpiration</b> | 7        | Specifies the number of days that obsolete copies of migrated files remain on the Tivoli Storage Manager server before they are expired. The range of values is 0 through 9999. See “ <b>migfileexpiration</b> ” on page 110.                                                                                                              |
| <b>migrateserver</b>     |          | 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 “ <b>migrateserver</b> ” on page 111.                                                                                                                                    |
| <b>minmigfilesize</b>    | 0        | Specifies the minimum file size for a file to be eligible for migration. The range of values is 0 through 2147483647. See “ <b>minmigfilesize</b> ” on page 111.                                                                                                                                                                           |
| <b>minrecalldaemons</b>  | 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 “ <b>minrecalldaemons</b> ” on page 112.                                                                                                                            |
| <b>optionformat</b>      | 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. Select standard or short format. See “ <b>optionformat</b> ” on page 112.                                                                                                    |
| <b>reconcileinterval</b> | 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 “ <b>reconcileinterval</b> ” on page 113.                                                                                                                                                |
| <b>restoremigstate</b>   | YES      | 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 YES or NO. See “ <b>restoremigstate</b> ” on page 113.                                                                                                   |
| <b>streamseq</b>         | 0        | Specifies the number of megabytes processed by the recall daemon before synchronizing the file contents. This option applies to streaming recall mode only. The range of values is 0 - 1024. See “ <b>streamseq</b> ” on page 115.                                                                                                         |

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 that you want to contact for space management services. For each **servername** entry, include a **commmethod** entry to specify the communication method to use for client and server communications.

- Issue a value for each option and remove the leading asterisk (\*). You can specify options for more than one server.

#### Related reference

“Tivoli Storage Manager for Space Management version 6.3 updates” on page xix

## Editing the options file `dsm.opt`

Most of the options in the `dsm.opt` file affect the backup-archive client. However, some options affect the 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 operates only on local file systems.

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

Table 12. HSM client options in the `dsm.opt` file

| Option                         | Default                                                                            | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------------------|------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>defaultserver</code>     | The server that you identify in the first stanza of the <code>dsm.sys</code> file. | Specifies the name of the default Tivoli Storage Manager server to contact for space management services when you do not specify a server name on the <code>migrateserver</code> option. The value of <code>defaultserver</code> in the <code>dsm.sys</code> file overrides <code>defaultserver</code> in the <code>dsm.opt</code> file. See “ <code>defaultserver</code> ” on page 98.                                                                              |
| <code>hsmgroupedmigrate</code> | NO                                                                                 | When this option is set to YES the HSM transaction grouping is activated. Multiple files are migrated in each transaction with the Tivoli Storage Manager server. The files are migrated when either the transaction byte limit or the transaction group limit is reached. The transaction group limit is specified with the <code>txngroupmax</code> option, which is set on the Tivoli Storage Manager server. See “ <code>hsmgroupedmigrate</code> ” on page 102. |
| <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 server for each client node. The value of <code>migrateserver</code> in the <code>dsm.sys</code> file overrides <code>migrateserver</code> in the <code>dsm.opt</code> file. See “ <code>migrateserver</code> ” on page 111.                                                                                                                  |
| <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. The file must have been backed up after migration or premigration. Tivoli Storage Manager records the migration state of files during backup, so only files that were migrated or premigrated before backup can be restored to stubbed (migrated) state. See “ <code>restoremigstate</code> ” on page 113.                                                                                                              |

- Set the `DSM_CONFIG` environment variable to point to your `dsm.opt` file.

#### Related tasks

“Setting environment variables” on page 33

---

## Optional setup features

You can specify which files are eligible for migration, exploit clustered file systems and LAN-free data transfer, start the HSM graphical user interface, and set logging options.

### Assigning management classes to files

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

The 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 *IBM Tivoli Storage Manager for UNIX and Linux 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 | None        | Specifies that a file is eligible for automatic and selective migration, selective migration only, or neither. The values for this attribute are:<br><br>Auto The file is eligible for both automatic and selective migration.<br>Selective The file is eligible for selective migration only.<br>None The file is not eligible for migration.<br><b>Note:</b> 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 spacemgtechnique attribute, files are not migrated from your workstation.                                                                                                                                         |
| automignonuse    | 0           | Specifies the number of days (0 through 9999) that must elapse since you last accessed the file before it is eligible for automatic migration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| migrequiresbkup  | YES         | Determines whether a current backup version of the file must exist on your migration server before the file is eligible for automatic or selective migration. The values for this attribute are:<br><br><b>YES</b> A current backup version must exist on the Tivoli Storage Manager server to which the file is migrated.<br><br><b>NO</b> A current backup version is not required.<br><b>Note:</b> 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 does not exist 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.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| 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"> <li>• Specify a different include-exclude options file for each server.</li> <li>• Assign only one management class to a file in a specific include-exclude options file.</li> <li>• Assign different management classes to files in different include-exclude options files.</li> </ul> <p>For example, if you back up files in /home/holland to Server1, the include-exclude options file that you use for Server1 might specify a management class named <i>mgmt1a</i> for a file named /home/holland/testfile. This management class must contain an appropriate backup copy group for the file.</p> <p>If you migrate files in the /home file system to Server2, the include-exclude options file that you use for that server might specify a management class named <i>mgmt2b</i> for the same file. That management class must contain appropriate space management attributes for the file.</p> |

## 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 following 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.

**Related reference**

“dsmmigquery” on page 146

## 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\_64 GPFS file systems do not use the **incl excl** option in a failover environment. Unlike `dsm.opt` and `dsm.sys` files, include-exclude options files are not shared between different nodes of a failover group. Add the include-exclude list directly to the `dsm.sys` file or verify that the include-exclude list files match on all nodes participating in the local failover group or node set.

For all other file systems, use the **incl excl** option in your `dsm.sys` file to specify the name of your include-exclude options file. You can create an include-exclude options file for each 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 Tivoli Storage Manager servers, you can include an **incl excl** option in each stanza. Each **incl excl** option can point to a different include-exclude options file. The files that you create must reside in a directory to which all users on your workstation have read access.

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

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

*Table 15. Include and Exclude Statements*

| Option                              | Description                                                                                                                                       |
|-------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>excl</b>                         | Excludes a file or a group of files from space management.                                                                                        |
| <code>exclude.backup</code>         | Excludes a file from backup.                                                                                                                      |
| <code>exclude.file</code>           | Excludes a file or a group of files from space management.                                                                                        |
| <code>exclude.file.spacemgmt</code> | Excludes a file from HSM services only. Use this option when you have both the backup-archive client and the HSM client installed.                |
| <code>exclude.spacemgmt</code>      | Excludes files and directories from HSM services only. Use this option when you have both the backup-archive client and the HSM client installed. |

Table 15. Include and Exclude Statements (continued)

| Option         | Description                                                           |
|----------------|-----------------------------------------------------------------------|
| <b>include</b> | Includes files for backup and HSM services or LAN-free data transfer. |
| include.file   | 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 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 **Add** to add a new include-exclude statement. You can select a statement that you want to update, remove, or move up or down the statements list.

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

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

## 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                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>DSM_DIR</b>    | Points to the resource files, the <code>dsm.sys</code> file, and the executable file, <code>dsmtca</code> . If you do not set <b>DSM_DIR</b> , the HSM client searches for the executable files in the installation directory.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>DSM_CONFIG</b> | Points to your <code>dsm.opt</code> file. <ul style="list-style-type: none"> <li>• If you do not set <b>DSM_CONFIG</b>, the HSM client searches for the options file in the directory to which <b>DSM_DIR</b> points.</li> <li>• If you do not set <b>DSM_DIR</b>, the HSM client searches for the options file in the installation directory.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>DSM_LOG</b>    | Points to the directory where you want the <code>dsmerror.log</code> 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"> <li>• If you define <b>DSM_DIR</b> but you do not define <b>DSM_LOG</b>, messages are written to <code>dsmerror.log</code> in the directory that you specified in <b>DSM_DIR</b>.</li> <li>• If you do not define <b>DSM_LOG</b> or <b>DSM_DIR</b>, error messages are written to <code>dsmerror.log</code> in the current directory. You receive a warning message if Tivoli Storage Manager cannot write messages to the log file. Processing continues.</li> </ul> |

## 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 the following is required:

- 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 Storage Agent User's Guide for AIX or Linux.

### LAN-free options

After installing and configuring the 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 set up LAN-free data transfer. For more information about these options, see *IBM Tivoli Storage Manager for UNIX and Linux Backup-Archive Clients Installation and User's Guide*.

Table 17. LAN-free data transfer options

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

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

**To include files:**

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

**To exclude files:**

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

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

## Setting up space management for HACMP on AIX JFS2

The HSM client permits your space-managed file system to become a part of a resource group in a HACMP environment 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 Tivoli Storage Manager Space Management clients support HACMP on JFS2 file systems at the following HACMP levels:

- HACMP 5.4.1 on AIX 6.1
- HACMP 5.5 on AIX 6.1

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.

## HSM client limitations for HACMP environments

You should be aware of the aspects of an HSM client HACMP environment that aren't completely integrated with the backup-archive client HACMP environment.

- The management class information is for the default migration server only.
- The server options information is for the default migration server only.
- Every HACMP computer must run the same HSM version.

**Note:** The HSM GUI does not support HACMP commands.

## Starting and stopping the HSM client in an HACMP environment

Use these scripts as examples that demonstrate how to start and stop the HSM client in an HACMP environment on an AIX file system.

**Note:** These scripts are only examples. You must create your own scripts for your environment.

Figure 1. *start\_HSM* : sample script to start HSM

```
#!/usr/bin/perl
#-----
call with
start_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 "* Starting HSM *" >> $LOG`;
`echo "*****" >> $LOG`;

#!/usr/bin/perl
#-----
call with
start_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 "* Starting HSM *" >> $LOG`;
`echo "*****" >> $LOG`;

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

killing all running daemons to reread the option files
my $pid = `ps -aef|awk '/dsmrecalld/{if (\$3 == "1") print \$2}'`;
`echo "Killing dsmrecalld ... $pid" >> $LOG`;
`kill -15 $pid 2>&1 1>> $LOG`;
`procwait -v $pid 2>&1 1>> $LOG`;

$pid = `ps -aef|awk '/dsmonitor{d/{if (\$3 == "1") print \$2}'`;
`echo "Killing dsmonitor{d ... $pid" >> $LOG`;
`kill -15 $pid 2>&1 1>> $LOG`;
`procwait -v $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`;
`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 '/dsmscoutd/{if (\ $3 == "1") print \ $2}'`;
`echo "Killing dsmscoutd ... $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\_64 GPFS clusters

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

The HSM client for AIX GPFS and Linux 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 are in the standard `/opt/tivoli` directory for Linux 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.

In an HACMP environment, you must configure HSM on every cluster node that manages or takes over space-managed file systems from other nodes.

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

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

You can control distributed recall by setting the option **hsmdistributedrecall** to YES on all source nodes that you do not want to recall files.

To enable distributed migration and recall capabilities follow these steps:

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

#### Related concepts

“Linux x86\_64 GPFS HSM client installation overview” on page 13

“AIX GPFS HSM client installation overview” on page 7

#### Related reference

“**hsmdistributedrecall**” on page 100

### Limitations of HSM support for AIX GPFS and Linux x86\_64 GPFS

Be aware of some limitations of HSM support for AIX General Parallel Files System (GPFS) and Linux x86\_64 GPFS file 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\_64 GPFS systems is not completely integrated with the backup-archive client support. For example, the HSM client refers to the *file\_system/.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\_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\_64 GPFS, 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 41 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
PortNumber	The PortNumber option specifies on which port the Space Management GUI can connect to the space management agent. Valid values are 1500 to 32152.
AgentTimeOut	The AgentTimeOut option should not be changed. You should only change the parameter if a service representative asks you to do so.
SessionTimeOut	The SessionTimeOut option should not be changed. You should only change the parameter if a service representative asks you to do so.
GuiOptionsFile	The GuiOptionsFile option is currently not used.
Tracing	The Tracing, TraceFile and TraceFlags option should only be used if a service representative asks you to do so.

## Setting HSM logging

The HSM log contains information about file migration and recall, threshold migration, reconciliation, and starting and stopping the HSM daemon.

You can analyze the HSM log to determine the current state of the system. For instance, the logs can indicate when a recall has started but not finished within the last hour. The administrator can analyze a particular recall and react accordingly.

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

Use the following options in the `dsm.sys` file to set up the log file:

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

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

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

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

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

| **Related reference**

| "Tivoli Storage Manager for Space Management version 6.3 updates" on page xix

| "**hsmlogeventflags**" on page 103

| "**hsmlogmax**" on page 103

| "**hsmlogname**" on page 104

| "**hsmlogretention**" on page 105

| "**hsmlogsampleinterval**" on page 106

| "dsmmigrate" on page 149

| "dsmrecall" on page 155

| "dsmmigundelete" on page 151

| "dsmdf" on page 126

| "**dsmmigfs query**" on page 140

---

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

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

### Note:

- To add and configure space management for a file system, you must have root user authority.
- The 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 does the following:

- Creates a hidden directory for the file system named `.SpaceMan` that stores certain information objects required to perform space management tasks.
- 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` file 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.
- Starts space management for the file system.

For AIX GPFS and Linux x86\_64 GPFS file systems, the `dsmswchd` daemon starts at system startup with an entry in the `etc/inittab` file or with the `initctl` service. Ensure that DMAPI is enabled on all GPFS file systems that the HSM client manages. Issue the following command to query this information:

```
/usr/lpp/mmfs/bin/mmfsfs DevicePath -z.
```

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

On GPFS 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 `dsmswchd` daemon is already set up on one of the cluster nodes within the GPFS cluster.

The AIX or Linux 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\_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\_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 do 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.
- You can add space management to a nested file system.
- You can add space management to an exported file system.

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

### Related tasks

“Adding space management to nested file systems” on page 45

“Adding space management to an exported file system” on page 45

### Related reference

Appendix A, “Control files in the .SpaceMan directory,” on page 171

---

## 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`.

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.

### Related reference

“`dsmmigfs add` and `update`” on page 132

---

## Adding space management to nested file systems

You can add space management to a nested file system.

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.

### Related tasks

“Adding space management to file systems” on page 44

## Mounting the parent file system before nested file system

For an AIX JFS2 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 the following command: `/usr/sbin/showmount -a`.
3. Add space management to your file system using the **dsmmigfs add** command.
4. To export the NFS file system again, issue the following command:  
`/usr/etc/exportfs -a .`
5. Instruct all NFS clients to mount the exported NFS file system again.

**Related reference**

“**dsmmigfs add** and **update**” on page 132

---

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

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

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

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

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

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

**Related reference**

“**dsmmigfs takeover**” on page 145

---

## Adding space management to WPARs on AIX 6.1 and 7.1

AIX 6.1 and 7.1 workload partitions (WPARs) act and look like 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 about editing the `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 are transparently recalled. They are transparently recalled if accessed from the global partition.

The HSM client is supported only in 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 .`

### Related reference

“**dsmmigfs add** and **update**” on page 132

---

## Updating settings from the HSM GUI

You can update space management settings to change the way the HSM client manages space on a 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.

---

## Settings that control the space usage of a file system

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

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

- The high and low thresholds for your file system that determine when threshold migration automatically starts and stops
- The total number of megabytes of data that you can migrate and premigrate from your file system to 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 qualify for partial file recall
- The minimum size (in megabytes) that a file must qualify for streaming recall mode
- The maximum number of files in the space-managed file system.

The information for each file system is stored in the `hsmfsconfig.xml` file, which is located in the appropriate `.SpaceMan` directory. You can modify the settings in the `hsmfsconfig.xml` file using the **`dsmmigfs update`** command. You can also use the **`dsmmigfs query`** command to display the settings in this file.

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

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

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

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

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

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

For information about configuring GPFS integration with the HSM client, see the Tivoli field guide *TSM for Space Management for UNIX-GPFS Integration Part I: Policy-driven Threshold Migration* at <https://www.ibm.com/support/docview.wss?uid=swg27018848>.

### Related concepts

Chapter 3, “Configuring the HSM client,” on page 23

### Related reference

“**dsmmigfs query**” on page 140

“**dsmmigfs add** and **update**” on page 132

“**dsmmigrate**” on page 149

## 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\_64 GPFS file systems, the file size must be greater than the stub file 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.

### Related reference

“**dsmmigfs add** and **update**” on page 132

“**minmigfilesize**” on page 111

## 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.

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

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

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

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

For information about configuring GPFS integration with the HSM client, see the Tivoli field guide *TSM for Space Management for UNIX-GPFS Integration Part I: Policy-driven Threshold Migration* at <https://www.ibm.com/support/docview.wss?uid=swg27018848>.

#### **Related reference**

**"dsmmigfs add and update"** on page 132

## **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.

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

For information about configuring GPFS integration with the HSM client, see the Tivoli field guide *TSM for Space Management for UNIX-GPFS Integration Part I: Policy-driven Threshold Migration* at <https://www.ibm.com/support/docview.wss?uid=swg27018848>.

#### **Related reference**

**"dsmmigfs add and update"** on page 132

## **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.

**Related reference**

“**dsmmigfs add** and **update**” on page 132

## 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. Leader data is especially useful if you frequently run programs that read only the information located at the beginning 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. Consider also the disadvantage of using additional space on your local file system for larger stub files.

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

For HSM clients on AIX Enhanced Journaled Files Systems (JFS2), 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, a smaller size 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, the operating system allocates one fragment for each stub file. The excess number of bytes remain unused. The default stub size is 4096.

For HSM clients on AIX GPFS and Linux x86\_64 GPFS file systems, you can specify 0 or a multiple of the file system fragment size. The default value is 0. A stub file on a GPFS file system occupies at least 8 KB because of the metadata in the stub file.

**Related reference**

“**dsmmigfs add** and **update**” on page 132

## Minimum stream file size

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

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

- AIX GPFS
- AIX JFS2
- Linux x86\_64 GPFS

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

#### Related reference

“**dsmmigfs add** and **update**” on page 132

### Minimum partial file size

In an AIX GPFS or Linux 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.

#### Related concepts

“Recalling migrated files overview” on page 3

Chapter 5, “Migrating files,” on page 61

### Maximum number of files

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

You can specify a maximum file size value 0 - 9999999999999999.

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

#### Related concepts

“Streaming recall mode” on page 79

#### Related reference

“**dsmmigfs add** and **update**” on page 132

“**dsmmigfs query**” on page 140

---

## Managing a file system with multiple Tivoli Storage Manager servers

### Overview

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

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

A file in a multiple-server environment is coupled with the Tivoli Storage Manager server to which it is initially migrated or backed up. The coupling between the file and the server is persistent over the life of both the file and the Tivoli Storage Manager server, and ensures data consistency. The file cannot be migrated or backed up to another Tivoli Storage Manager server until the file is uncoupled. Automatic migration and transparent recall processes contact the server with which a file is coupled.

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

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

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

#### **Related tasks**

“Restoring a file system in an environment that is managed by multiple Tivoli Storage Manager servers” on page 74

#### **Related reference**

“Tivoli Storage Manager for Space Management version 6.3 updates” on page xix

## **Enabling a file system to be managed by multiple Tivoli Storage Manager servers**

Enable multiple-server support for all HSM clients in the GPFSnode, then add Tivoli Storage Manager servers to manage the file system.

In a multiple-server environment, the Tivoli Storage Manager server to which a file is backed up must be the server to which a file is migrated. Before the upgrade, each file system on a GPFS cluster that is managed by an HSM client must be managed by a single Tivoli Storage Manager server. That means that files on a file system must be backed up and migrated to the same Tivoli Storage Manager server.

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

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

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

For information about configuring GPFS integration with the HSM client, see the Tivoli field guide *TSM for Space Management for UNIX-GPFS Integration Part I: Policy-driven Threshold Migration* at <https://www.ibm.com/support/docview.wss?uid=swg27018848>.

For information about using the GPFS backup command **mmbackup** with a file system that is space-managed by multiple Tivoli Storage Manager servers, see the GPFS FAQs topic in the IBM Cluster information center: [http://publib.boulder.ibm.com/infocenter/clresctr/vxrx/index.jsp?topic=/com.ibm.cluster.gpfs.doc/gpfs\\_faqs/gpfs\\_faqs.html](http://publib.boulder.ibm.com/infocenter/clresctr/vxrx/index.jsp?topic=/com.ibm.cluster.gpfs.doc/gpfs_faqs/gpfs_faqs.html).

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

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

To query status of the multiple-server environment, issue the following command: **dsmmigfs querymultiserver file\_system\_name**. You can add more Tivoli Storage Manager servers to the list of servers that manage a file system.

#### Related reference

“**dsmmigfs addmultiserver**, **querymultiserver**, and **removemultiserver**” on page 136  
 “**hsmmultiserver**” on page 107  
 “**hsmdisableautomigdaemons**” on page 99

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

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

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

1. Add space management to the file system with the following command: **dsmmigfs add -server=server\_name file\_system\_name options**.
2. Add support for multiple servers with the following command: **dsmmigfs addmultiserver -server=server\_name file\_system\_name**. You must specify the same Tivoli Storage Manager server in the **server** option in steps 1 and 2.

The HSM client now manages space on the file system.

To query status of the multiple-server environment issue the following command: **dsmmigfs querymultiserver file\_system\_name**. You can add more Tivoli Storage Manager servers to the list of servers that manage the file system.

## Related reference

“**dsmmigfs add** and **update**” on page 132

“**dsmmigfs addmultiserver**, **querymultiserver**, and **removemultiserver**” on page 136

## Removing a Tivoli Storage Manager server from a multiple-server environment

You can remove a Tivoli Storage Manager server from an environment that supports multiple servers.

Before you remove a server from managing a file system, recall all files from the Tivoli Storage Manager server to the local file system. Then remove the server with the following steps:

1. Remove the Tivoli Storage Manager server from the multiple server list by issuing the following command : **dsmmigfs removemultiserver** **-server=server\_name file\_system\_name**. Run the command for all file systems that were managed by the server. After running the command, files are not backed up or migrated to the server that was removed. You can still recall migrated copies and restore backup copies from the server that was removed.
2. Run the `dsmMultiServerRemove.pl` script.

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

- a. Recalls all migrated files that are coupled with the Tivoli Storage Manager server that was removed to the specified file system.  
The HSM client determines how much space is required to recall all files to the file system. If there is not enough space, the HSM client notifies you. To remove the server, make space available and run the script again.
- b. Reconciles the specified file system with the Tivoli Storage Manager server that was removed. The following reconciliation tasks are performed:
  - The HSM client verifies that valid objects for each local stub file exist in the space management pool of the specified Tivoli Storage Manager server.
  - When orphans are identified, their names are recorded in the `.SpaceMan/orphan.stubs` file.
  - The HSM client expires all files in the space management storage pool on the Tivoli Storage Manager server.

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

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

- c. Uncouples all files that are coupled with the Tivoli Storage Manager server that was removed.

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

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

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

After files are decoupled from the Tivoli Storage Manager server that was removed, they can be migrated and backed up to another server.

If at least one other Tivoli Storage Manager server manages the file system, the file system is still space-managed.

#### **Related reference**

“`dsmmigfs addmultiserver`, `querymultiserver`, and `removemultiserver`” on page 136  
“`dsmRemoveServer.pl`” on page 169

## **Limitations for multiple-server environments**

Some features are not supported when a file system is managed by multiple Tivoli Storage Manager servers.

### **Tivoli Storage Manager server name encryption**

It is not possible to encrypt multiple Tivoli Storage Manager server names on a single node.

### **Tivoli Storage Manager server node replication**

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

### **Tivoli Storage Manager server LAN-free configuration**

A LAN-free Tivoli Storage Manager server configuration is not supported in a multiple-server environment.

## **Replicating nodes in a multiple-server environment**

If a Tivoli Storage Manager becomes unavailable, the node replication feature provides a temporary replacement server. In an HSM multiple-server environment, you must take several steps to switch from the source server to a target server.

The Tivoli Storage Manager node replication feature provides disaster recovery capabilities and reduces the down time of a productive backup system. A node replication environment consists of two Tivoli Storage Manager servers: a source server and a target server. The source server incrementally replicates all data and metadata, ensuring the completeness and consistency of the data and metadata. When the source server becomes unavailable, the backup-archive client or the HSM client can switch to the target server. The target server can be used only for restore, retrieve and recall processes. The target server cannot be used for backup, archive, or migration processes. Tivoli Storage Manager node replication replaces the productive backup server temporarily for maintenance or in the event of an unplanned outage.

For more information about the node replication feature, see "Replicating client node data" in IBM Tivoli Storage Manager Administrator's Guide.

In a multiple-server environment, each Tivoli Storage Manager server includes storage pools for backup and migration. It is important to ensure that all pools used for backup and migration are replicated to the same target server.

To switch from a source server to a target server in a multiple server environment do the following steps:

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

You can recall migrated files from the target server. You cannot back up or migrate files to the target server. To back up or migrate files, you must switch back to the source server. To switch back to the source server, repeat the steps above, but change the value of the **servername** option from the target server to the source server.

---

## 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 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 before deactivation.
- 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. You can also use 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. The client node must contact the new server for space management services. The administrator can use a **lock node** command to prevent the client node from migrating or recalling files before 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:

```
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 must 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**.

### Related reference

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

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

---

## Reactivating space management

If you deactivated space management from your file systems, you can reactivate space management 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
```

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. You can also use 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 must 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. If the node was globally deactivated, click **Global Reactivate** from the menu. Otherwise, to activate a deactivated file system from the work area, select the file system, and click the **Reactivate** button.

### Related reference

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

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

---

## Removing space management

You can remove space management from a file system.

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. Ensure also 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 does 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.
- 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
```

### 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 must 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 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 are recalled from the server.

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

#### Related tasks

“Resolving orphaned stub files” on page 87

#### Related reference

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

---

## Chapter 5. Migrating files

The HSM client automatically migrates files to Tivoli Storage Manager storage when space is needed on the local file system. On the local system, stub files replace these migrated files. The system administrator can control 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 the last access date of the file to change. This change influences the file processing of the HSM client, since the last access date is one of the factors that determines when a file becomes eligible for migration.

The Tivoli Storage Manager backup-archive client uses option `preserveLastAccessDate` to determine whether to reset the last access date of any specified files following a backup or archive operation. By default, the backup-archive client does not reset the last access dates to their original values 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 1 byte and the file system block size.
- It meets management class criteria.
- A file in a multiple-server environment that has been coupled with a server can only be migrated to that server.

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

In a GPFS environment, newly created files that are migrated either automatically or selectively must be older than 2 minutes before they can be migrated. Migrating files less than 5 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 the files. The `dsmdf` command will display correct results after GPFS synchronization and after the next reconciliation of the file system.

### Related concepts

Chapter 1, “HSM client for UNIX and Linux overview,” on page 1

“Migrating files overview” on page 2

### Related reference

“dsmdf” on page 126

“dsmdu” on page 127

---

## 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 are premigrated as specified with the option **pmpercentage** parameter of the **dsmmigfs** command.

Demand migration responds to an out-of-space condition on your local file system. The space monitor daemon checks for an out-of-space condition every 10 seconds. Threshold migration starts automatically if the used capacity of your file system exceeds a certain limit. The default for this limit is 90 percent. For example, if you attempt to copy a large file into your file system, and there is not enough available space for the file, eligible files begin migrating automatically from your local file system to 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\_64 GPFS which stops the process). You only get an out-of-space error message on AIX GPFS and Linux 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.

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

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

For information about configuring GPFS integration with the HSM client, see the Tivoli field guide *TSM for Space Management for UNIX-GPFS Integration Part I: Policy-driven Threshold Migration* at <https://www.ibm.com/support/docview.wss?uid=swg27018848>.

#### Related reference

“**dsmmigfs add** and **update**” on page 132

---

## Candidate selection for automatic migration

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

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

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

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

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

The CFI is updated by various processes:

- The daemon periodically scans the managed file systems and updates the CFI with the latest information about every file. The information is used to generate a list of files that is used for threshold and demand migration. Before migration, each file is checked to determine 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.

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

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

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

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

For information about configuring GPFS integration with the HSM client, see the Tivoli field guide *TSM for Space Management for UNIX-GPFS Integration Part I: Policy-driven Threshold Migration* at <https://www.ibm.com/support/docview.wss?uid=swg27018848>.

**Related concepts**

“The scout daemon” on page 90

**Related reference**

“**minmigfilesize**” on page 111

“**maxcandprocs**” on page 108

---

## File premigration

For faster migration, the HSM client prepares files for automatic migration using a process called *premigration*.

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.

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

Use the **dsmautomig** command to reduce space usage to the low threshold on your file system before it reaches the high threshold. using the . If space usage exceeds the low threshold 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. Running automatic migration on one pool until it reaches the low threshold does not result in a low threshold for the entire file system. To reach low threshold for the entire file system, issue the **dsmautomig** command for the entire file system (without a storage pool argument). Low threshold is enforced for each storage pool in that file system.

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

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

```
dsmautomig /home
```

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

### Related reference

“dsmautomig” on page 124

---

## 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 are not using some 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 for the file does not change. Unlike automatic migration, the number of days since you last accessed a file has no effect on whether your file is eligible for selective migration.

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

An eligible file must meet the following management class requirements:

- The management class assigned to the file enables 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 :

- 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 started from a shell application or a script to build a list of files to be recalled. 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
```

#### Related reference

“dsmmigrate” on page 149

“hsmgroupedmigrate” on page 102

---

## Migrating a single file system to two or more Tivoli Storage Manager servers

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

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

#### Related concepts

“Managing a file system with multiple Tivoli Storage Manager servers” on page 53

---

## Chapter 6. Backing up and restoring file systems

Coordinate backup and migration to protect the file system data.

When you back up a file with the backup-archive client, a copy of the file is created on the 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 a 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.

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 back up 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 space-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 only need to 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, ensure 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, results in 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 migrating them. Use the backup-archive client to restore the actual backup versions only.

#### **Related concepts**

Chapter 1, “HSM client for UNIX and Linux overview,” on page 1

---

## **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:** 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. This behavior applies to the following file systems:

- AIX GPFS
- Linux x86\_64 GPFS

### **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 **restoremigstate** option in your `dsm.opt` file. The file is restored to your local file system as a resident file.

**Note:** Files with ACLs are restored to a resident state, even when you specify YES on the **restoremigstate** option in your `dsm.opt` file. This behavior applies to the following file systems:

- AIX GPFS
- Linux x86\_64 GPFS

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 restoring so that you fully understand your range of options.

Each step shows you a different method for restoring files.

1. If you back up and migrate files to the same server, you can restore a backup version of the file to the local file system. 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.
2. If you back up and migrate files to the same server, you can restore a stub file to the local file system. Set the **restoremigstate** option to YES in your `dsm.opt` file.

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

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

### Related concepts

"Recreating stub files using the **dsmmigundelete** command"

### Related reference

"**restoremigstate**" on page 113

"**migfileexpiration**" on page 110

---

## Recreating stub files using the **dsmmigundelete** command

The **dsmmigundelete** command uses the migrated files on the Tivoli Storage Manager server to re-create stub files.

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

### Note:

- You cannot use the **dsmmigundelete** command to 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 backup versions of your stub files are available on a Tivoli Storage Manager server, use the backup-archive client to restore your stub files.
- The **dsmmigundelete** command does not support hard linked files. If you attempt to recreate a stub file for a hard linked file, a stub file is not recreated unless all of the files that are hard linked together are deleted from your local file system. When one file in a set of hard linked files is migrated, all of the hard linked files in the set become stub files. When the **dsmmigundelete** command recreates a stub file for a hard linked file, the stub file has the same name as the

file that was originally migrated. Stub files are not recreated for any other files that were previously in the hard linked set of files.

The following actions occur when you issue the **dsmmigundele** command:

- The **dsmmigundele** 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 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 **dsmmigundele** 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. dsmmigundele 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 <b>dsmmigundele</b> command is finished.
Premigrated files	Creates a stub file for a premigrated file if a corresponding original file does not exist on your local file system, and the premigrated file was not marked for expiration. The premigrated file will be unexpired at the server after the <b>dsmmigundele</b> command is finished.

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

*Table 20. dsmmigundele 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.

Table 20. `dsmmigundelete` command actions with the `expiring` option (continued)

File	Description
Premigrated 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.

#### Related reference

"`dsmmigundelete`" on page 151

---

## 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 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. The HSM client restores backup versions of resident files but restores migrated and premigrated files to stub files.

#### Related tasks

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

"Restoring a file system backed up and migrated to a different server" on page 73

## 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.
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.

For more information about the backup-archive client, see *IBM Tivoli Storage Manager for UNIX and Linux Backup-Archive Clients Installation and User's Guide*.

#### Related tasks

"Adding space management to file systems" on page 44

#### Related reference

"**dsmmigfs add** and **update**" on page 132

## 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 **dsmc restore** command, see *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.
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/*"
```

#### Related concepts

“Recreating stub files using the **dsmmigundelete** command” on page 70

#### Related tasks

“Adding space management to file systems” on page 44

#### Related reference

“**dsmmigfs add** and **update**” on page 132

“**dsmmigundelete**” on page 151

## Restoring a file system in an environment that is managed by multiple Tivoli Storage Manager servers

You can recover files that were migrated and backed up from a single file system to multiple Tivoli Storage Manager servers.

This task assumes that you migrated and backed up from a single file system to multiple Tivoli Storage Manager servers.

1. Determine the Tivoli Storage Manager servers that manage the file system. Issue the following command: **dsmmigfs querymultiserver** *file\_system*.
2. Restore the directory structure. For each Tivoli Storage Manager server that manages the file system, issue the following command: **dsmc restore -dironly -latest -server=migration\_server file\_system\_name**.
3. Restore the stub files. For each Tivoli Storage Manager server that manages the file system, issue the following command: **dsmmigundelete -server=migration\_server file\_system**. The **dsmmigundelete** command creates stub files and couples the stub files with the Tivoli Storage Manager server.
4. Restore files that were not migrated or premigrated. For each Tivoli Storage Manager server that manages the file system, issue the following command: **dsmc restore -ifnewer -server=migration\_server file\_system**. The **ifnewer** option ensures that the latest version of a file is recovered.

#### Related concepts

“Managing a file system with multiple Tivoli Storage Manager servers” on page 53

#### Related reference

“**dsmmigfs addmultiserver**, **querymultiserver**, and **removemultiserver**” on page 136

“**dsmmigundelete**” on page 151

---

## 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 the Tivoli Storage Manager server. To minimize the impact of such losses, create a set of installation media that can restore your system to a state that permits contact with the Tivoli Storage Manager 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 Tivoli Storage Manager backup-archive clients, you must reinstall the clients before you can recover your files.

The installation media must 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.

After you restore your operating system and are ready to restore your file system, do one of the following tasks. If you back up files and migrate files to the same server, see “Restoring a file system backed up and migrated to the same server” on page 72. If you back up and migrate files to different servers see “Restoring a file system backed up and migrated to a different server” on page 73.

### Related tasks

“Restoring a file system backed up and migrated to the same server” on page 72

“Restoring a file system backed up and migrated to a different server” on page 73



---

## 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 known 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.

### Related concepts

Chapter 1, “HSM client for UNIX and Linux overview,” on page 1

“Recalling migrated files overview” on page 3

### Related reference

“dsmattr” on page 122

---

## Transparent recall processing

Transparent recall processing 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 processing

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

For example, if some of your files are stored on a disk storage device and some of your files are stored on a tape storage device, the 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
```

#### Related concepts

“Recalling migrated files overview” on page 3

#### Related reference

“dsmrecall” on page 155

---

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

When a vendor application makes a read request for a file that is qualified for partial file recall, and the file is migrated, the 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 **minpartialrecallsize** option value. Use the **dsmrecall** command with the **offset** and **size** options to specify the data range within the file.

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

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

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

- Set the **minpartialrecallsize** option of the **dsmmigfs** command to the minimum size (in megabytes) that a file must have to qualify for partial file recall.

- Set the **recallmode** option of the **dsmattr** command to `partialrecall`. This specifies that, regardless of its size, a file is recalled using partial file recall. Use this method to change the recall mode of migrated files that you normally read but you do not modify. When you set the recall mode to partial file recall, this mode remains associated with your files until you:
  - Change the recall mode
  - Modify the file
  - Recall the file selectively
  - Restore the file

#### Related concepts

“Recalling migrated files overview” on page 3

#### Related reference

“**dsmmigfs add** and **update**” on page 132

“**dsmattr**” on page 122

“**dsmrecall**” on page 155

## Streaming recall mode

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

#### Note:

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

You can use the following methods to specify which files 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
- Set the **minstreamfilesize** option of the **dsmmigfs update** command to specify a number to enable or disable an asynchronous recall of migrated files.

A larger value of the **streamseq** option increases the performance of streaming recall.

### Related concepts

“Recalling migrated files overview” on page 3

### Related reference

“**streamseq**” on page 115

“**dsmattr**” on page 122

“**dsmmigfs add** and **update**” on page 132

---

## 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.

### Related reference

“**dsmattr**” on page 122

---

## Optimized tape recall processing

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

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

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

Optimized tape recall processing provides the following benefits:

- Avoids frequent tape mount and unmount operations.
- Avoids excessive tape seek operations.
- Enables tape drives to go to streaming mode, if possible, for optimal performance. (Streaming mode is possible only if the recalled files are located contiguously on the tape).
- Orders and recalls files in separate steps so that you can restrict the recall operations on certain tapes.
- Can recall files from several tape drives in parallel to increase recall throughput.

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

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

### Contention for tape resources

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

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

#### Related concepts

“Recalling migrated files overview” on page 3

#### Related reference

“Tivoli Storage Manager for Space Management version 6.3 updates” on page xix  
 “dsmrecall” on page 155

“**hsmmaxrecalltapedrives**” on page 106

## List files for optimized tape recalls

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

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

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

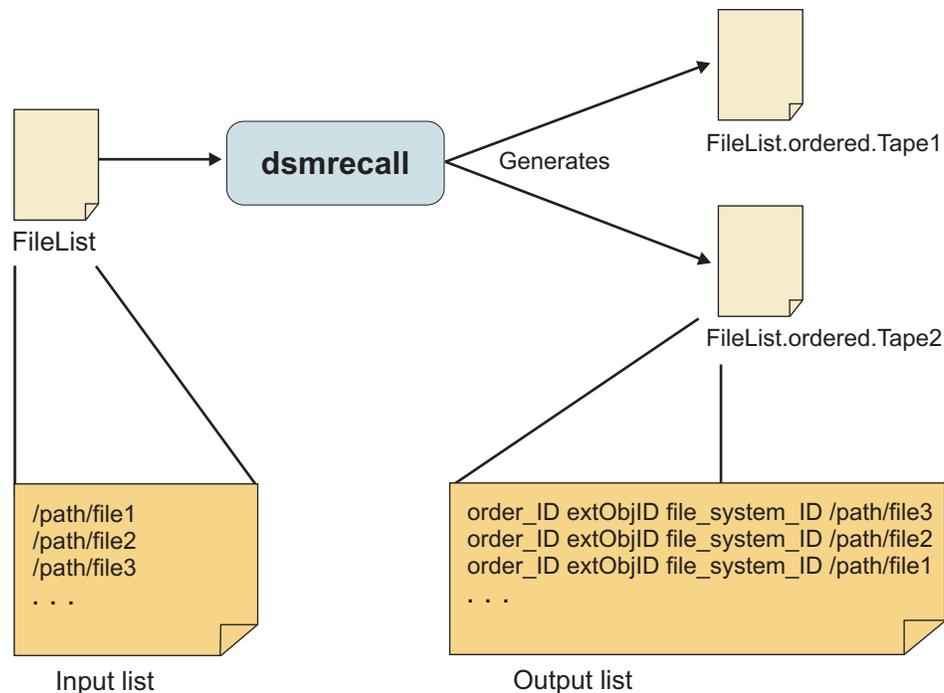


Figure 4. Ordered list file generation

## Input file lists

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

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

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

## Generated list files

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

## Tape and disk list files

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

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

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

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

## Collection files

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

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

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

- Entries in the collection file have this format:

```
filecount byteCount filelist.nonTape.00000000
filecount byteCount filelist.ordered.tape tapeID1
filecount byteCount filelist.ordered.tape tapeID2
```

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

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

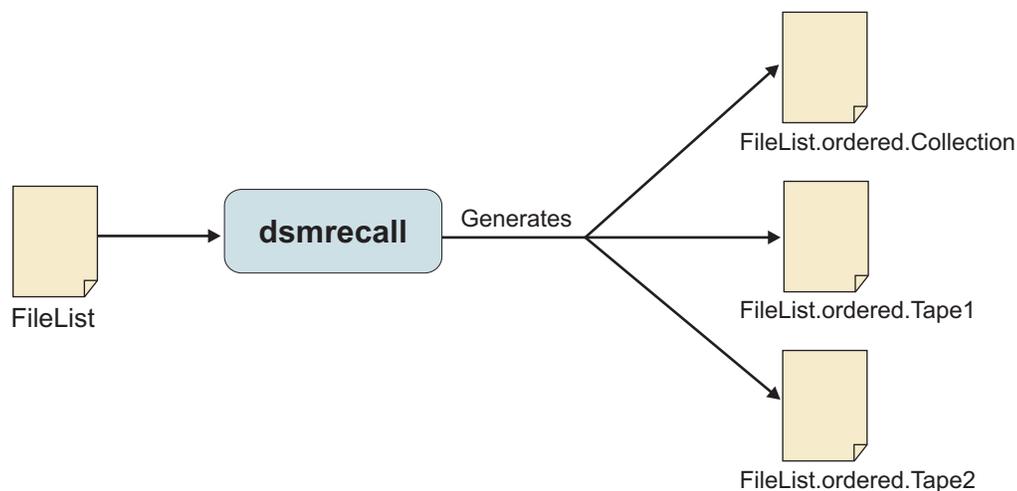


Figure 5. Generation of list files

| **Related concepts**

| Chapter 11, "HSM client dsm.sys and dsm.opt option reference," on page 95

| **Related reference**

| "dsmrecall" on page 155

| **Recalling files with optimization for tape**

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

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

| Start a selective recall with the **filelist=list\_file** option and specify a file system. For example:

| `dsmrecall -filelist=myFileList myFileSystem`

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

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

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

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

| **Related concepts**

| "List files for optimized tape recalls" on page 81

| **Related reference**

| "dsmrecall" on page 155

---

## Chapter 8. Reconciling file systems

To keep local file systems synchronized with the Tivoli Storage Manager server for space management services, the HSM client automatically reconciles your file systems at preset intervals. 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.

### Related concepts

Chapter 1, “HSM client for UNIX and Linux overview,” on page 1

“Reconciling file systems overview” on page 4

---

### 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 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 95.

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"><li>• Verifies that a stub file exists on your local file system for each migrated file in storage</li><li>• Marks a migrated file for expiration if:<ul style="list-style-type: none"><li>– You deleted the stub file from your local file system</li><li>– You recalled the file and modified it</li></ul></li><li>• Removes a migrated file from storage if it expired</li><li>• Updates the status file</li></ul>

Table 22. Automatic reconciliation tasks (continued)

File type	Reconciliation Tasks
Premigrated files	<ul style="list-style-type: none"><li>• Verifies that premigrated files are still valid</li><li>• Marks a premigrated file for expiration if:<ul style="list-style-type: none"><li>– You deleted the file from your local file system</li><li>– You accessed the file and modified it</li></ul></li><li>• Removes a copy of a premigrated file from storage if it expired.</li><li>• Updates the status file</li></ul>
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 87 for more information.
Status file	<p>Updates the following information in the status file:</p> <ul style="list-style-type: none"><li>• Number of premigrated files</li><li>• Number of premigrated blocks</li><li>• Number of migrated files</li><li>• Number of migrated blocks</li></ul> <p><b>Note:</b> 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 reconciliation 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 158 and “The scout daemon” on page 90 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.”

For command information see “dsmreconcile” on page 158.

---

## 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.

---

## Reconciliation based on a GPFS policy

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

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

| The HSM GPFS client can use the GPFS policy engine to reconcile file systems  
| with Tivoli Storage Manager servers. In a multiple-server environment, you must  
| specify the Tivoli Storage Manager server to reconcile.

| To reconcile a file system, issue the following command: `dsmreconcileGPFS.pl`  
| *file\_system\_name*.

| In a multiple-server environment you must specify a server to reconcile with the  
| **server** option: `dsmreconcileGPFS.pl -server=server_name file_system_name`. To  
| reconcile the file system with all servers, start the `dsmreconcileGPFS.pl` script with  
| each Tivoli Storage Manager server.

| The `dsmreconcileGPFS.pl` script uses the GPFS policy engine to determine which  
| files are migrated. The script also starts the **dsmreconcile** command to perform a  
| two-way reconciliation by using the **orphancheck** and **fileinfo** options.

| **Related reference**

| “`dsmreconcileGPFS.pl`” on page 168

| “`dsmreconcile`” on page 158

---

## Chapter 9. Space management daemons

The HSM client uses the space monitor, recall, scout, watch, and root daemons to manage file systems automatically.

These daemons are included with the Tivoli Storage Manager HSM client. The space monitor daemon (**dsmonitor**), the scout daemon (**dsmscoutd**), the recall daemon (**dsrecalld**), and the root daemon (**dsrootd**) manage space on JFS2 and GPFS file systems. The watch daemon (**dswatchd**) runs only on GPFS file systems.

The daemons start when you add space management to your file systems and when you modify space-management options.

On GPFS file systems, you can use the GPFS policy engine to monitor space thresholds, search for migration candidates, and identify migrated files for reconciliation. If you use the GPFS policy engine, the space monitor and scout daemons can be disabled.

For information about configuring GPFS integration with the HSM client, see the Tivoli field guide *TSM for Space Management for UNIX-GPFS Integration Part I: Policy-driven Threshold Migration* at <https://www.ibm.com/support/docview.wss?uid=swg27018848>.

---

### The space monitor daemon

The space monitor daemon monitors space usage on all file systems and the storage pools in those file systems to which you add space management. It starts threshold migration whenever necessary.

To check space usage more frequently or less frequently, change the value on the **checkthresholds** option in the `dsm.sys` file. To reconcile your file systems more frequently or less frequently, change the value on the **reconcileinterval** option in the `dsm.sys` file. When you change the option values the space monitor daemon uses, the new values are not effective until you stop and restart the space monitor daemon.

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

When GPFS policy drives automatic migration, disable the space monitor daemon by setting the option **hsmdisableautomigdaemons=YES**.

For information about configuring GPFS integration with the HSM client, see the Tivoli field guide *TSM for Space Management for UNIX-GPFS Integration Part I: Policy-driven Threshold Migration* at <https://www.ibm.com/support/docview.wss?uid=swg27018848>.

#### Related reference

“**reconcileinterval**” on page 113

“**checkthresholds**” on page 96

“**dsmmonitord**” on page 153

“**hsmdisableautomigdaemons**” on page 99

---

## The recall daemon

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

If a recall daemon is not running, enter the **dsmrecalld** command to start one.

A child recall daemon can recall only one file at a time; however, you can run more than one recall daemon at the same time. To set the minimum and maximum number of child recall daemons that you want to run at one time, use the **minrecalldaemons** and **maxrecalldaemons** options in the `dsm.sys` file. The minimum number of child recall daemons that you can run at the same time is 1. The default is 3. The maximum number of child daemons that you can run at the same time is 99. The default is 20. On a General Parallel File System (GPFS) the maximum is constrained by the GPFS option **dmapiWorkerThreads**.

If all child recall daemons are busy another file cannot be recalled until a child recall daemon is available. If an application uses all available child recall daemons because all files are migrated, the application is suspended until a child recall daemon is available. You can increase the value of the **maxrecalldaemons** option.

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

#### Related reference

“**maxrecalldaemons**” on page 109

“**minrecalldaemons**” on page 112

“**dsmwatchd**” on page 164

“**dsmrecalld**” on page 158

---

## The scout daemon

The scout daemon automatically searches for candidates on each file system or storage pool in a file system for which space management is active.

The scout daemon scans file systems and stores the information for each file in a complete file index (CFI). The daemon works with the CFI to search for migration candidates. The CFI is updated automatically during all migration, recall, and restore operations.

To specify how often the scout daemon scans a file system, modify the setting on the **candidatesinterval** option in the `dsm.sys` file.

When GPFS policy drives automatic migration, disable the scout daemon by setting the option **hsmdisableautomigdaemons=YES**.

| For information about configuring GPFS integration with the HSM client, see the  
| Tivoli field guide *TSM for Space Management for UNIX-GPFS Integration Part I:  
| Policy-driven Threshold Migration* at [https://www.ibm.com/support/  
| docview.wss?uid=swg27018848](https://www.ibm.com/support/docview.wss?uid=swg27018848).

**Related reference**

“dsmscoutd” on page 162

“candidatesinterval” on page 95

---

## The watch daemon

If failover processing is active on your General Parallel Files System (GPFS) node, the watch daemon checks the status of other daemons. If any of the other daemons end or become corrupted, the watch daemon automatically recovers the failed daemon.

The watch daemon checks the status of the recall, the monitor, the scout, and the root daemons. If the watch daemon must recover the root daemon but cannot start it, ensure that the port mapper service is running.

| If two or more GPFS nodes participate in a failover, the watch daemon can take  
| over the file systems of a failed HSM node.

| 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\_64 GPFS only. It replaces the  
| Reliable Scalable Cluster Technology (RSCT) Group Services functions for the HSM  
| client. Therefore, you do not need to create an RSCT peer domain for the HSM  
| client.

| On Red Hat Enterprise Linux Version 6 (RHEL6), the **dsmwatchd** daemon is started  
| from the **initctl** service. It can be started manually with the following command:  
| `initctl start HSM`. Stop the watch daemon with the following command: `initctl  
| stop HSM`.

On all other systems besides RHEL6, you cannot start the watch daemon manually. It is started by the **init** service.

**Related reference**

“dsmwatchd” on page 164

---

## 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.

If the root daemon is not running, enter the **dsmrootd** command to start it. It is important that the port mapping service is running.

**Related reference**

“dsmrootd” on page 161

---

## Stopping the space management daemons

Use the following steps to properly stop space management daemons.

Do not use the **kill -9** command to stop any space management daemon. All daemons have their own clean-up procedure, which is interrupted if you use the **kill -9** command. Using the **kill -9** command can yield unpredictable and unintentional results.

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

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

1. Issue the **dsmq** command to obtain the recall ID and the recall daemon process ID for each recall process that is in the queue.
2. Issue the **dsmrm** command to remove each recall process from the queue.
3. Issue **ps -ef | grep dsm** to verify that both the space monitor daemon and the master recall daemon are running.
4. Issue the **kill -15** command with the process identifier number to stop the daemons.
5. For AIX and Linux x86\_64 GPFS file systems, issue the command **dmkilld** to stop the recall daemons.
6. Verify that the daemons are no longer running. For AIX and Linux GPFS file systems, issue the **dsmmigfs stop** command to stop all space management daemons.
7. Stop the **dsmwatchd** daemon.

**On Red Hat Enterprise Linux Version 6 (RHEL6) systems:**

Stop the daemon with the following command: `initctl stop HSM.`

**On all other systems besides RHEL6:**

- a. Comment out the entry for the **dsmwatchd** daemon in the file `/etc/inittab`.
- b. Close the file and submit the change by issuing the following command: `telinit Q`.
- c. Stop the **dsmwatchd** daemon with the command `kill -15`.

**Related reference**

“**dsmq**” on page 154

“**dsmrm**” on page 161

“**dmkilld**” on page 121

“**dsmmigfs stop, start, and restart**” on page 144

---

## 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.

**Related reference**

"**reconcileinterval**" on page 113

---

## 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.

### Related tasks

“Editing the options file `dsm.opt`” on page 28

“Editing the options file `dsm.sys`” on page 25

---

### `candidatesinterval`

The `candidatesinterval` option specifies how often the `dsmscoutd` daemon searches the file systems for migration candidates.

The `dsmscoutd` daemon scans the file systems and stores the information in a complete file index (CFI), which is used to search for migration candidates.

**Tip:** When automatic migration is driven by GPFS policy, the `dsmscoutd` daemon is disabled, and the `candidatesinterval` option has no effect.

### Options File

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

### Syntax

►—`CANDIDATESINTERVAL—scan_interval`—►

### Parameters

#### *scan\_interval*

Specifies the maximum interval of time that elapses between each successive time the `dsmscoutd` daemon automatically scans for files in the file systems. The range of values is 0 to 9999. The default is 1.

Specify 0 to continuously scan the file systems. When scanning reaches the end of a file system, the `dsmscoutd` daemon immediately begins scanning again from the beginning of the file system.

Specify 1 to scan the file systems at intervals that depend on the percentage of the file system content that has changed. The `dsmscoutd` daemon increases the frequency of scanning as the percentage of file system changes increases. The `dsmscoutd` daemon reduces the frequency of scanning as the percentage of file system changes decreases. This is the default.

Specify a value from 2 to 9999 to define the number of hours between file system scans. After scanning reaches the end of the file system, the `dsmscoutd` daemon waits the specified number of hours before starting the next scan.

### Related reference

“`dsmscoutd`” on page 162

---

## checkfororphans

The **checkfororphans** option specifies whether or not the **dsmreconcile** command checks for migrated files that are no longer present on the server but whose stub files still remain on the client (orphans). The option parameter that you set determines whether or not the **dsmreconcile** command queries the scout daemon for all migrated and premigrated files.

### Options File

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

### Syntax



### 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.

### Related reference

“**dsmreconcile**” on page 158

“**dsm Scoutd**” on page 162

---

## checkthresholds

The **checkthresholds** option specifies how often the space monitor daemon checks space usage on your file systems. The space monitor daemon checks each file system to which you added space management.

**Tip:** When automatic migration is driven by GPFS policy, the `dsmonitord` daemon is disabled, and the **checkthresholds** option has no effect.

### Options File

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

### Syntax



## Parameters

### *interval*

Specifies the number of minutes that must elapse before the space monitor daemon checks space usage on your file systems. The range of values is 1 through 9999. The default is 5.

### Related reference

“dsmmonitord” on page 153

---

## 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.

**Tip:** This option controls compression *only if* your administrator specifies that your client node determines the selection. The server also can define this option.

A fast processor on a slow network connection benefits from compression, but a slow processor on a fast network connection does not. Use this option with the backup-archive client option, **compressalways**.

If you specify **compressalways YES**, compression continues even if the file size increases. To stop compression if the file size grows, and resend the uncompressed file, specify **compressalways NO**.

If you specify **compression YES**, you can control compression processing in the following ways:

- Use the **exclude.compression** option in your include-exclude options file to exclude specific files or groups of files from compression processing.
- Use the **include.compression** option in your include-exclude options file to include files within a broad group of excluded files for compression processing.

## Options File

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

## Syntax



## Parameters

**NO** Files are not compressed before they are sent to the server. This is the default.

### YES

Files are compressed before they are sent to the server.

---

## defaultserver

The **defaultserver** option specifies the default server to which you back up and archive your files from your local file systems. If you do not specify a migration server with the **migrateserver** option, this option can also specify the server to which files are migrated from your local file systems.

### Options File

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

### Syntax

►—DEFAULTServer—*servername*—◄

### Parameters

*servername*

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

### Related reference

“**migrateserver**” on page 111

---

## errorlogname

The **errorlogname** option specifies the path of the file that logs error messages for the HSM client and the backup-archive client.

This option specifies the fully qualified path and file name of the error log file. If this option is not specified, the default log file is used. The default error log file is the `dsmerror.log` file in the current working directory.

Authorized user: The value for this option overrides the **DSM\_LOG** environment variable.

The log file path cannot be a symbolic link.

### Options File

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

## Syntax

►—ERRORLOGName—*filespec*—►

## Parameters

*filespec*

The fully qualified path in which to store error log information. If any part of the path you specify does not exist, the HSM client creates it.

The log file path cannot be a symbolic link.

---

## errorprog

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

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

The program is started only in case of very severe errors such as DMI or file system errors. Also, only processes that do not have an output channel, such as daemons, use this option. Use the full path to specify the program. The error message is sent to the standard input of the program.

## Options File

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

## Syntax

►—ERRORProg—*program-name*—►

## Parameters

*program-name*

Specifies the path and file name of the program to which you want to send a message if a severe error occurs during space management processing.

Some examples:

```
errorprog /usr/bin/cat >/tmp/tsm_severe_errors.txt
```

```
errorprog /usr/bin/perl /root/dsmsevererror.pl
```

---

## hsmdisableautomigdaemons

Use the **hsmdisableautomigdaemons** option to control the start of the **dsmscouthd** and **dsmmonitord** daemons.

## Supported Clients

This option is valid for only AIX GPFS clients and Linux x86\_64 GPFS clients.

When a GPFS policy monitors migration thresholds, the **dsmscouthd** and **dsmmonitord** daemons can be disabled. Stop and restart the **dsmwatchd** daemon to

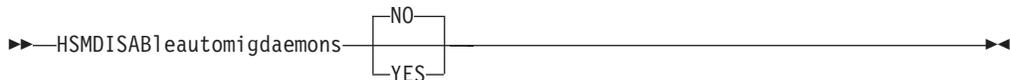
activate a new value of the **hsmdisableautomigdaemons** option. The **dsmwatchd** daemon is restarted by the **init** process or the **initctl** process.

Valid for only AIX GPFS clients and Linux x86\_64 GPFS clients.

## Options File

Place this option in the client options file **dsm.opt**.

## Syntax



## Parameters

**NO** Specifies that the HSM client automatic migration daemons **dsm scoutd** and **dsmmonitord** are started. This is the default.

### YES

Specifies that the **dsm scoutd** and **dsmmonitord** daemons are not started.

**Tip:** The **dsm scoutd** and **dsmmonitord** daemons can also be started by the GPFS daemon (**mmfsd**).

## Related tasks

“Stopping the space management daemons” on page 92

---

## hsmdistributedrecall

The **hsmdistributedrecall** option specifies whether a node that is running as owner node can delegate recall requests to other nodes in a GPFS cluster.

To stop distributed recalls for an entire cluster, set this option to **NO** for all HSM nodes in the cluster and restart the **dsmrecalld** daemons. Then, only the HSM owner node can recall migrated files to the space-managed file systems it owns.

## Options File

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

## Syntax



## Parameters

### YES

The node can delegate recall requests to other nodes in a GPFS cluster. This is the default.

**NO** The node cannot delegate recall requests to other nodes in a GPFS cluster.

---

## hsmenableimmediatemigrate

The **hsmenableimmediatemigrate** option determines if files that are less than 2 minutes old can be migrated during selective migration.

When a file is created, file status information can take some time to become synchronized among GPFS nodes. If you create a file, then migrate the file before file information is synchronized, the file size information returned from a **dsmdu** or **dsmdf** command can be inaccurate. Because of the typical time needed for GPFS synchronization, by default the HSM client does not migrate a file that is less than 2 minutes old.

To migrate files that are less than 2 minutes old, set **hsmenableimmediatemigrate=YES**. When you set **hsmenableimmediatemigrate=YES**, the HSM client issues system synchronization commands before migrating each file that is less than 2 minutes old.

This option applies only during selective migration.

### Options File

Place this option in the client options file `dsm.opt`.

### Syntax



### Parameters

**NO** Files that are less than 2 minutes old are not migrated. This is the default.

#### YES

Files that are less than 2 minutes old can be migrated during selective migration.

#### Related reference

“`dsmmigrate`” on page 149

---

## hsmextobjidattr

The **hsmextobjidattr** option specifies whether the **dsmreconcile**, **dsmmigrate**, **dsmrecall**, or **dsmrecalld** command assigns an object ID to the file as a Data Management Application Programming Interface (DMAPI) extended attribute (`extObjId`).

The `extObjId` DMAPI extended attribute is required for two-way orphan check reconciliation.

The two-way orphan check processing identifies orphans on the file system and on the Tivoli Storage Manager server. The two-way check is done in a single pass and the orphan identification process uses parallel processing.

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

## Syntax



## Parameters

**NO** An object ID is not assigned to the file that is processed by the **dsmmigrate**, **dsmrecall**, **dsmrecalld**, or the **dsmreconcile** command with the **preptwo** option. When a reconciliation is run with the **preptwo** option, files are not prepared for the two-way orphan check. This is the default.

### YES

An object ID is stored in the file that is processed by the **dsmmigrate**, **dsmrecall**, **dsmrecalld**, or the **dsmreconcile** command with the **preptwo** option.

---

## hsmgroupedmigrate

The **hsmgroupedmigrate** option specifies whether the **dsmmigrate** command migrates more than one file with each transaction.

The transfer of multiple files per transaction can improve performance when migrating many small or medium size files (up to 100 MB).

The default action is to migrate each file in a single transaction and make a stub file when the transaction is successful.

When this option is set to YES the HSM transaction grouping is activated. A group of files is migrated in each transaction with the Tivoli Storage Manager server when either the transaction byte limit or the transaction group limit is reached. The transaction byte limit has a default of 25 MB, and can be set with the **txnbytelimit** option. The transaction group limit is specified with the **txngroupmax** option. The **txngroupmax** option is set on the server in the options file or on the node definition on the server. If a file with a different management class setting is added to the group, the group is migrated immediately.

If a file with a different file system root is added to the group, the group is migrated immediately.

## Options File

Place this option in the client options file `dsm.opt`.

## Syntax



## Parameters

**NO** The **dsmmigrate** command migrates one file per transaction with the Tivoli Storage Manager server.

YES

The **dsmmigrate** command migrates a group of files in each transaction with the Tivoli Storage Manager server. This is the default.

---

## hsmlogeventflags

The **hsmlogeventflags** option specifies which kinds of events are added to the HSM log file.

You must use an authorized user ID to use this option.

If this option is not specified, or if no flags are included in the specification, no HSM log entries are made.

### Options File

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

### Syntax

▶—HSMLOGEVENTFLAGS—*Flags*—▶

### Parameters

*Flags*

Indicate one or more of the following values:

**FILE**

File events are logged. File events include file migration and recall processing.

**FS** File system events are logged. File system events include threshold migration, reconciliation, and HSM-related file system configuration changes.

**SYSTEM**

System events are logged. System events include starting and stopping the HSM daemon.

---

## hsmlogmax

The **hsmlogmax** option specifies the maximum size of the HSM log, in megabytes.

You must use an authorized user ID to use this option.

If you use the **hsmlogretention** option instead of the **hsmlogmax** option, all existing log entries are retained and the log is pruned according to the new **hsmlogretention** option age criteria.

If you use the **hsmlogmax** option instead of the **hsmlogretention** option, all records in the existing log are copied to the pruned log file, `dsmhsm.pru`. Then the existing log is emptied, and logging begins under the **hsmlogmax** option size criteria.

If you change the value of the **hsmlogmax** option, the existing log is extended or shortened to accommodate the new size. If the value is reduced, the oldest entries are deleted to reduce the file to the new size.

**Restriction:** You cannot specify a nonzero **hsmlogmax** option value and activate the **hsmlogretention** option.

## Options File

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

## Syntax

▶▶—HSMLOGMAX—*size*—————▶▶

## Parameters

*size*

Specifies the maximum size, in megabytes, for the log file. The range of values is 0 - 2047; the default is 0, which specifies that the log file has no size limit.

## Related reference

“**hsmlogretention**” on page 105

---

## hsmlogname

The **hsmlogname** option specifies the HSM log file name.

This option specifies the fully qualified path and file name of the HSM log file. If this option is not specified, the default log file is used. The default log file is `dsmhsm.log` and is in the same directory as the `dsmerror.log` file.

You must use an authorized user ID to use this option.

The value for this option overrides the **DSM\_LOG** environment variable.

## Options File

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

## Syntax

▶▶—HSMLOGName—*filespec*—————▶▶

## Parameters

*filespec*

The fully qualified path and file name of the HSM log file. If any part of the path you specify does not exist, the HSM client creates it.

The log file path cannot be a symbolic link.

---

## hsmlogretention

The **hsmlogretention** option specifies how many days to maintain HSM log entries before pruning, and whether to save the pruned entries.

The HSM log is pruned when the first entry is written to the log after a Tivoli Storage Manager for Space Management command or daemon is started. Note that some HSM client daemons run continuously. Stop the daemons and start them again or activate one of the HSM commands to prune the HSM log.

You must use an authorized user ID to use this option.

If you use the **hsmlogmax** option instead of the **hsmlogretention** option, all records in the existing log are copied to the pruned log file, `dsmhsm.pru`. Then the existing log is emptied, and logging begins under the new size criteria.

If you use the **hsmlogretention** option instead of the **hsmlogmax** option, all existing log entries are retained and the log is pruned according to the new **hsmlogretention** option age criteria.

**Restriction:** You cannot specify the **hsmlogretention** option and a nonzero value for the **hsmlogmax** option.

### Options File

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

### Syntax



### Parameters

#### N or *days*

Specifies how long to wait before pruning the HSM log.

N Do not prune the log. The HSM log grows indefinitely. This is the default.

#### *days*

The number of days to keep log file entries before pruning the log. The range of values is 0 - 9999.

#### D or S

Specifies whether to save the pruned entries. Enter a space or comma to separate this parameter from the previous one.

D Discard the HSM log entries when you prune the log. This is the default.

S Save the HSM log entries when you prune the log.

The pruned entries are copied from the HSM log to the `dsmhsm.pru` file that is located in the same directory as the HSM log.

| **Related reference**

| "hsmlogmax" on page 103

| **Related information**

| Chapter 12, "HSM client command reference," on page 117

---

| **hsmlogsampleinterval**

| The **hsmlogsampleinterval** option specifies how often to gather file system data for the HSM log file.

| You must be an authorized user to use the **hsmlogsampleinterval** option.

| **Options File**

| Place this option in the client system options file `dsm.sys` within a server stanza.

| **Syntax**

| **▶▶**—HSMLOGSampleinterval—*interval*—**▶▶**

| **Parameters**

| *interval*

| Specifies the number of seconds between each gathering of file system data.  
| The range of values is 0 - 9999999 and the default value is 3600 (one hour). If 0  
| is specified, no file system data is gathered.

---

| **hsmmaxrecalltapedrives**

| The **hsmmaxrecalltapedrives** option controls how many tape drives participate in tape-optimized recall processing.

| By limiting the available tape drives for tape optimized recall processing, you can prevent all drives from being blocked. In this way, some drives remain accessible for transparent recall processing and migration.

| If the value of the **hsmmaxrecalltapedrives** option is greater than the number of tape drives available, a recall thread waits for the next available tape drive. If you plan to run tape-optimized recalls processes in parallel, the number can be decreased accordingly for parallel processing without running out of tape drives.

| If the **maxnummp** option is defined on the Tivoli Storage Manager server, you can set the **hsmmaxrecalltapedrives** option to a number that is less than or equal to the **maxnummp** option. You can reserve tape drives for other operations if you set the number lower than the **maxnummp** option and the maximum number of tape drives.

| **Options File**

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

| **Syntax**

| **▶▶**—HSMMAXREcalltapedrives—*value*—**▶▶**

## Parameters

### *value*

Specifies the number of tape drives available for a tape-optimized recall. The range of values is 1 -10. The default is 5.

---

## hsmmultiserver

The **hsmmultiserver** option enables and disables migration and backup in an environment of multiple Tivoli Storage Manager servers.

## Supported Clients

You can enable migration and backup from a single file system to multiple Tivoli Storage Manager servers. You must set **hsmmultiserver=YES** for all HSM clients in the GPFS cluster.

If you set **hsmmultiserver=NO**, all of the command options for the multiple-server environment are disabled. When you set **hsmmultiserver=NO**, you can recall files only from the migration server that is configured in the `dsm.sys` options file. If a file was migrated to a Tivoli Storage Manager server other than the migration server that is configured in the `dsm.sys` options file, you cannot recall the file.

The option is available only for AIX GPFS clients and Linux GPFS clients.

## Options File

Place this option in the client options file `dsm.opt`.

## Syntax



## Parameters

**NO** All of the command options for the multiple-server environment are disabled. All HSM commands have standard input and output. This is the default.

### **YES**

All of the command options for the multiple-server environment are enabled. All HSM commands that display a list add an additional column that shows the Tivoli Storage Manager server name.

## Related information

Chapter 12, “HSM client command reference,” on page 117

---

## incl excl

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

### Syntax

▶—incl excl—*filespec*—————▶

### Parameters

*filespec*

Specifies the path and file name for your include-exclude file.

### Related concepts

“Modifying the include-exclude options file” on page 31

### Related tasks

“Creating an include-exclude list” on page 32

---

## maxcandprocs

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

**Tip:** When automatic migration is driven by GPFS policy, the **dsmscoutd** daemon is disabled, and the **maxcandprocs** option has no effect.

### Options File

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

### Syntax

▶—MAXCANDProcs—*number*—————▶

### Parameters

*number*

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

---

## maxmigrators

The **maxmigrators** option specifies the maximum number of parallel migration sessions 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.

**Tip:** When automatic migration is driven by GPFS policy, the **maxmigrators** option has no effect.

Do not set the **maxmigrators** option higher than the number of parallel sessions that the server can use to store data.

This option can be set by the Tivoli Storage Manager server.

## Options File

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

## Syntax

▶▶—MAXMIGRators—*number*—▶▶

## Parameters

*number*

Specifies the maximum number of parallel migration sessions that you can set. The range is 1 - 20. The default is 5. If this option is changed from the default, 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 the 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.

**Note:** The GPFS configuration option **dmapiWorkerThreads** can limit the maximum concurrent recall tasks to less than the value of **maxrecalldaemons**.

## Options File

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

## Syntax

▶▶—MAXRECALLdaemons—*number*—▶▶

## Parameters

*number*

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

---

## maxthresholdproc

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

**Tip:** When automatic migration is driven by GPFS policy, the **maxthresholdproc** option has no effect.

### Options File

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

### Syntax

▶▶—MAXThresholdproc—*number*—————▶▶

### Parameters

*number*

Specifies the maximum number of automatic threshold migration processes that the HSM client can start at one time. The range of values is 1 - 99. The default is 3.

---

## migfileexpiration

The **migfileexpiration** option specifies the number of days that copies of migrated or premigrated files remain on the server after they are modified on your local file system or are deleted from your local file system.

### Options File

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

### Syntax

▶▶—MIGFileexpiration—*days*—————▶▶

### Parameters

*days*

Specifies the number of days a copy of a migrated or premigrated file remains in storage after it is modified on your local file system, or deleted from your local file system. The range of values is 0 - 9999. The default is 7 days.

**Note:** If you specify a value of 0, an obsolete copy of a migrated or premigrated file is deleted from the server during the next reconciliation run. If you delete a file from the local file system and reconcile runs with the `migfileexpiration` value as 0, the file can not be recreated by the **dsmmigundelete** process.

---

## migrateserver

The **migrateserver** option specifies the name of the server to which you want to migrate files from your client node. Specify one migration server for each client node.

If you do not specify a server with the **migrateserver** option, your files migrate to the server that you specify with the **defaultserver** option. If you do not specify a server with either of these options, your files migrate to the server that you identify in the first stanza of your `dsm.sys` file.

You can override this option with the following command: `dsmmigfs upd /FS -Server=servername`. Replace `servername` with the name of your server.

The value of **migrateserver** in `dsm.sys` overrides **migrateserver** in `dsm.opt`.

After your files migrate to the server that you specified, do not specify a different migration server unless your administrator transfers your migrated files from the specified server to another. Otherwise, the server cannot locate your migrated files until you specify the server to which your files were originally migrated.

### Options File

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

### Syntax

►—MIGRateserver—*servername*—◄

### Parameters

*servername*

Specifies the name of the server to which you want to migrate files from your client node. Your `dsm.sys` file must contain a stanza beginning with the **servername** option and it must contain the required communication options for the server that you specify with the **migrateserver** option.

### Related reference

“**defaultserver**” on page 98

---

## minmigfilesize

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

This option applies to all of the space-managed file systems for which you have not specified a file system specific value for **minmigfilesize** using the **dsmmigfs add** or **dsmmigfs add** commands. See “**dsmmigfs add** and **update**” on page 132 for details.

### Options File

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

## Syntax

►►—MINMIGfilesize—*fileSize*—◄◄

## Parameters

*fileSize*

Specifies the minimum file size, in bytes, for a file to be eligible for migration. The range of values is 0 through 2147483647. The default is 0.

For AIX GPFS and Linux 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\_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.

## Options File

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

## Syntax

►►—MINRecalldaemons—*number*—◄◄

## Parameters

*number*

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

---

## optionformat

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

The `optionformat` option is a deprecated option. It is supported in version 6.3, but will not be supported in future versions. In future versions, all options must be specified with the standard format.

## Options File

Place this option in the `dsm.opt` or `dsm.sys` 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.

**Tip:** When **hsmdisableautomigdaemons=YES**, the **reconcileinterval** option has no effect.

## Options File

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

## Syntax

►►—RECOncileinterval—*interval*—►►

## Parameters

*interval*

Specifies the number of hours that must elapse between each successive time your file systems are automatically reconciled on your workstation. If you specify a value of 0, your file systems are not reconciled automatically. The range of values is 0 - 9999. The default is 24.

---

## restoremigstate

The **restoremigstate** option specifies whether you want to restore or retrieve stub files or backup-archive versions of migrated files during a restore-retrieve operation. Use this option with the backup-archive client **restore** and **retrieve** commands.

Place this option in the `dsm.opt` or `dsm.sys` file.

You can restore or retrieve a stub file for a migrated file only when the following is true:

- The file exists in the migration storage pool
- The file is backed up or archived and migrated to the same server

When the number of days elapse that you specified with the **migfileexpiration** option, the migrated file is removed from storage.

If you specify **restoremigstate** YES, and if the migrated file has not expired, the file is restored or retrieved to a stub file, regardless of whether it is marked for expiration.

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

The **restoremigstate** option restores a file if it is backed up after migration. If the file is backed up *before* migration, you cannot restore a stub file because a server stub file copy does not exist.

Files with access control lists (ACLs) are restored in a resident state regardless of the setting for **restoremigstate**. 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.

## Options File

| Place this option in the dsm.opt or dsm.sys file.

### Syntax



## Parameters

### YES

Restores or retrieves migrated files to stub files on your local file system during a restore or retrieve operation. The files remain migrated. This is the default.

**Note:** A stub file 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.

---

## streamseq

The **streamseq** option specifies the number of megabytes processed by the recall daemon before synchronizing the file contents. This option applies to streaming recall mode only.

### Options File

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

### Syntax

▶—STReamseq—*megabytes*—▶

### Parameters

#### *megabytes*

Specifies the maximum number of megabytes that the recall daemon processes. The range of values is 0 - 1024. The default is 0.

The recall daemon writes the file contents in asynchronous mode until the specified value is reached. The file contents are then synchronized. Another synchronization occurs at the end of the file stream.

When the value is 0, the recall daemon synchronizes the file content at the end of each block written. This value might have a negative impact on performance for a large-file recall operations. If system performance is affected, specify a value for the **streamseq** option that is not 0. For optimal performance, set the value of the **streamseq** option to a value that is a multiple of the preferred block size of the file system. Use the following commands to determine the preferred block size of the file system:

#### For AIX JFS2:

**Command:** `lsjfs2 /home`

#### Output:

```
#MountPoint:Device:Vfs:Nodename:Type:Size:Options:
AutoMount:Acct:OtherOptions:LvSize:FsSize:BlockSize:
Sparse:InlineLog:InlineLogSz:EAformat:Quota:DMAPI:VIX:

/home:/dev/hd1:jfs2:::32768:rw:yes:no::32768:32768:
4096:yes:no:0:v1:no:no:no:
```

#### For AIX/Linux GPFS

**Command:** `mmlsfs /dev/lvgpfs03 -B`

#### Output:

flag	value	description
----	-----	-----
-B	262144	Block size

**Important:** If you use the **streamseq** option to set the preferred block size of a file system, this can impact recalls to another file system having a different preferred block size.



---

## 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.

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.
- Do not stop any HSM process with the command `kill -9`.
- Use the following wildcard characters in file, directory, or file system specifications. The shell in which you are running matches and expands wildcard characters.
  - \* Matches zero or more characters
  - ? Matches any single character
- 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.

### Related concepts

Chapter 3, “Configuring the HSM client,” on page 23

---

## Standard and short option formats

You can specify options on some commands. These options are specified differently dependent upon the value of the **optionformat** option.

### Note:

The **optionformat** option is a deprecated option. It is supported in version 6.3, but will not be supported in future versions. In future versions, all options must be specified with the standard format.

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 -mgmtclass -detail /home
```
- Separate each option with a blank space.
- Issue more than one option in a command in any order before or after a file, directory, or file system specification.
- If the option defines a value, separate the option name from the value with an equal sign (=). For example:

```
dsmmigfs update -ht=90 /home
```

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 -g -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 -gv /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	<pre>dsmmigrate -recursive -detail /home/user1/file1 dsmmigrate -rec -det          /home/user1/file1 dsmmigrate -r -d              /home/user1/file1 dsmmigrate /home/user1/file1 -r -d</pre>
optionformat=short	<pre>dsmmigrate -R -v /home/user1/file1</pre>

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	<pre>dsmmigfs update -hthreshold=90 -lthreshold=20 /home dsmmigfs update -hthresh=90   -lthreshold=20 /home dsmmigfs update -ht=90 -l=20 /home dsmmigfs update /home -ht=90 -l=20</pre>
optionformat=short	<pre>dsmmigfs update -t90 -l20 /home dsmmigfs update -t90l20 /home dsmmigfs update /home -t90l20</pre>

---

## 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` 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
<b>dsmdf</b>	<p>Displays space usage information for a file system. For example, to display space usage information for the /home file system, issue the following command:</p> <pre>dsmdf /home</pre> <p>See “dsmdf” on page 126 for more information about this command.</p>
<b>dsmls</b>	<p>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:</p> <pre>dsmls /home/user1/*</pre> <p>See “dsmls” on page 129 for more information about this command.</p>
<b>dsmdu</b>	<p>Displays space usage information for files and directories. For example, to display space usage information for each file in the /home/user/proj1 directory and in all of its subdirectories, issue the following command:</p> <pre>dsmdu -Allfiles /home/user1/proj1</pre> <p>See “dsmdu” on page 127 for more information about this command.</p>
<b>dsmmigfs query</b>	<p>Displays the current space management settings for a file system. For example, to display the space management settings for the /home file system, issue the following command:</p> <pre>dsmmigfs query /home</pre> <p>See “dsmmigfs query” on page 140 for more information about this command.</p>
<b>dsmmigundelete</b>	<p>Recreates deleted stub files for migrated files and creates stub files for premigrated files if a corresponding original file does not exist on your local file system. The file then becomes a migrated file. For example, to recreate stub files for migrated files in the /home file system that are not marked for expiration (reconciliation was not run since the files were deleted) issue the command:</p> <pre>dsmmigundelete /home</pre> <p>See “dsmmigundelete” on page 151 for more information about this command.</p>

---

## 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
<b>dmkilld</b>	Only valid on AIX GPFS and Linux x86_64 GPFS  Stops the master recall daemon and all of its children, and interrupts all active recalls. See “dmkilld” on page 121.
<b>dsmattr</b>	Only valid for: AIX JFS2, AIX GPFS, Linux x86_64 GPFS  Sets or displays the recall mode for a migrated file. See “dsmattr” on page 122.
<b>dsmautomig</b>	Starts parallel migration sessions for a file system. See “dsmautomig” on page 124.
<b>dsmdf</b>	Displays space usage information for a file system. See “dsmdf” on page 126.
<b>dsmdu</b>	Displays space usage information for files and directories. See “dsmdu” on page 127.
<b>dsmls</b>	Lists files in a directory and displays file state. See “dsmls” on page 129.
<b>dsmmigfs add, update</b>	Adds space management to a file system, or updates space management attributes for a file system. See “ <b>dsmmigfs add</b> and <b>update</b> ” on page 132.
<b>dsmmigfs addmultiserver, querymultiserver, removemultiserver</b>	Adds space management to a file system, or updates space management attributes for a file system. See “ <b>dsmmigfs addmultiserver, querymultiserver, and removemultiserver</b> ” on page 136.
<b>dsmmigfs deactivate, reactivate, remove</b>	Deactivates or reactivates space management for a file system, or removes space management from a file system. See “ <b>dsmmigfs deactivate, reactivate, and remove</b> ” on page 137.
<b>dsmmigfs sdrreset, enablefailover, disablefailover</b>	Manages recovery from partial system failure (GPFS only). See “ <b>dsmmigfs sdrreset, enablefailover, and disablefailover</b> ” on page 143.
<b>dsmmigfs query</b>	Displays current space management settings for a file system. See “ <b>dsmmigfs query</b> ” on page 140.
<b>dsmmigfs globaldeactivate, globalreactivate</b>	Deactivates or reactivates space management for a space-managed client node. See “ <b>dsmmigfs globaldeactivate and globalreactivate</b> ” on page 138.
<b>dsmmigfs rollback</b>	Only valid for AIX GPFS and Linux 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 “ <b>dsmmigfs rollback</b> ” on page 142.
<b>dsmmigfs stop, start, restart</b>	Starts or starts the HSM daemons. See “ <b>dsmmigfs stop, start, and restart</b> ” on page 144.
<b>dsmmigfs takeover</b>	Transfers the HSM management of a file system to a HSM client node within the same local GPFS cluster. See “ <b>dsmmigfs takeover</b> ” on page 145.
<b>dsmmighelp</b>	Displays online help for commands. See “dsmmighelp” on page 146.
<b>dsmmigquery</b>	Displays space management information.  There are many space management and backup-archive client shared options available using the <b>dsmmigquery -o</b> command. See “HSM and backup-archive client dsmmigquery command shared options” on page 148 for a list of these shared options. See “dsmmigquery” on page 146.

Table 26. HSM command summary (continued)

Command and location	Description
<b>dsmmigrate</b>	Moves selected files from your local file system to Tivoli Storage Manager storage. See “dsmmigrate” on page 149.
<b>dsmmigundelete</b>	Recreates deleted stub files. See “dsmmigundelete” on page 151.
<b>dsmmonitord</b>	Starts the space monitor daemon. See “dsmmonitord” on page 153.
<b>dsmq</b>	Displays information, including recall IDs, for all files that are currently queued for recall. See “dsmq” on page 154.
<b>dsmrecall</b>	Moves selected files from storage to your local file system. See “dsmrecall” on page 155.
<b>dsmrecalld</b>	Starts the recall daemon. See “dsmrecalld” on page 158.
<b>dsmreconcile</b>	Synchronizes the client and server. See “dsmreconcile” on page 158.
<b>dsmrm</b>	Removes a recall process from the recall queue. See “dsmrm” on page 161.
<b>dsmrootd</b>	Starts the root daemon. See “dsmrootd” on page 161.
<b>dsmscoutd</b>	Starts, stops and restarts the scout daemon and shows file system information. See “dsmscoutd” on page 162.
<b>dsmsetpw</b>	Changes the Tivoli Storage Manager password for your client node. See “dsmsetpw” on page 163.
<b>dsmwatchd</b>	Only valid on AIX GPFS and Linux x86_64 GPFS  Manages failover activities for your HSM client node. See “dsmwatchd” on page 164.

## 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\_64 GPFS only.

### Syntax

```

|
| >>—DMKILLD—[options]
|

```

### Parameters

*options*

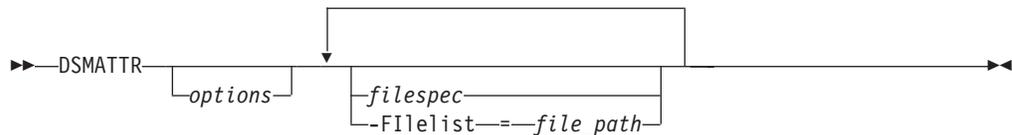
**-help** or when **optionformat=short: -h**

Displays syntax and options for the command. Do not specify other options when you specify the **help** option.

## dsmattr

The **dsmattr** command sets or changes the recall mode for one or more migrated files. The recall mode determines how the 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.

### Syntax



### Parameters

#### *options*

##### **-ERRORLOGName=***file\_path*

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file *dsm.sys* or *dsm.opt*, or as specified by the **DSM\_LOG** environment variable. If any part of the path you specify does not exist, the HSM client creates it.

##### **-Help** or when **optionformat=short**: **-h**

Displays syntax and options for the command. Do not specify other options when you specify the **help** option.

##### **-RECALLmode=***value* or when **optionformat=short**: **-rvalue**

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.
normal	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 is modified, it becomes a resident file.
partialrecall (AIX GPFS, Linux x86_64 GPFS only)	Specifies that the file should be recalled using partial file recall, regardless of its size.
Streaming	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 when **optionformat=short**: **-R**

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

#### *filespec*

**Note:** This parameter is required only when you set a new recall mode. If you do not use the **RECALLmode** option, and you do not specify a path and a file name, the current recall mode displays for all files in the current directory.

The path and file name of the file for which you want to set a new recall mode, or display the current recall mode. You can specify a single file, a group of files, or a directory. If you specify a directory, the 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=file\_path** or when **optionformat=short: -fi=file\_path**

Specifies the path of a file that contains a list of files to be processed by this command.

The entries in the list file must adhere to the following rules:

- Each entry is a fully qualified path to a file or a relative path to a file.
- No entry is a directory object.
- Each entry is on a separate line.

The HSM client ignores any entry that does not adhere to these rules.

The following rules apply to a file list for the HSM client. These rules differ from the rules for backup-archive client file lists.

- An entry can contain unprintable characters, but cannot contain a carriage return.
- Wildcard characters are allowed.
- A file name that contains spaces does not have to be enclosed in quotation marks.
- An entry that begins with a quotation mark *and* ends with a quotation mark is tolerated. The HSM client assumes that the quotation marks are not needed and omits them when processing the entry. If beginning and ending quotation marks are required to identify the file, add double quotation marks before the beginning and after the end of the entry.

The following is an example of a list of files within a list file:

```
/home/dir/file1
"/fs1/dir2/file3"
"/fs2/my files/file4"
/fs2/my files/file5
../tivoli/'file1'
"'fs3'/dir3/'file.txt'"
fs4/dir/a"file".txt
'/fs4/dir/file.txt'
/fs5/dir/file*with?wildcards.txt
```

If the list file name that you specify with the **filelist** option does not exist, the command fails.

## Examples

Task	Command
Change the recall mode to partial file recall for all migrated files in the /home/user2 directory and all of its subdirectories.	<code>dsmattr -recall=partialrecall -Recursive /home/user2</code>
Enable an asynchronous recall of migrated files in the /home/user2/ directory.	<code>dsmattr -recall=streaming /home/user2/</code>

Task	Command
Display the recall modes that are assigned to all files in the current directory.	<code>dsmattr</code>
Display all files in the filelist named <code>/tmp/filelist</code> .	<code>dsmattr -filelist=/tmp/filelist</code>

### Related concepts

Chapter 7, “Recalling migrated files,” on page 77

## 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 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.

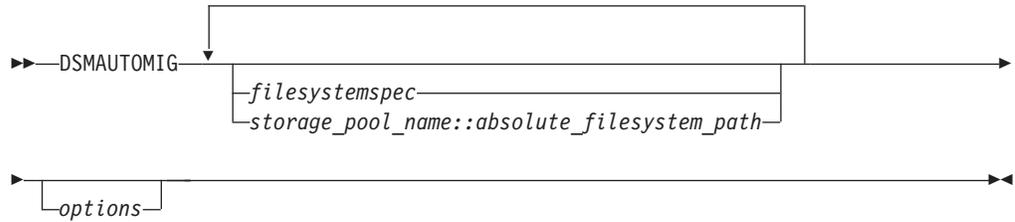
Hidden directories and files are included in automatic migration. These can be excluded from automatic migration by adding the hidden directories or files to the exclude list in the `dsm.opt` file.

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

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

**Note:** The Tivoli Storage Manager for Space Management client does not migrate contents of symbolic links.

## Syntax



## Parameters

### *filesystemspec*

Specifies the name of the file system for which you want to run threshold migration. The default is all file systems for which space management is active. You can specify more than one file system name, and you can use wildcard characters within a file system name. If you specify more than one file system name, separate each name with one or more blank spaces.

### *storage\_pool\_name::absolute\_filesystem\_path*

Specifies the storage pools that are located in the *absolute\_filesystem\_path* that are to be migrated automatically.

### *options*

#### **-Detail or when optionformat=short: -v**

Displays information about migrated files.

#### **-ERRORLOGName=file\_path**

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file *dsm.sys* or *dsm.opt*, or as specified by the **DSM\_LOG** environment variable. If any part of the path you specify does not exist, the HSM client creates it.

#### **-Help or when optionformat=short: -h**

Displays syntax and options for the command. Do not specify other options when you specify the **help** option.

## Examples

Task	Command
Start threshold migration for all storage pools on all file systems for which space management is active.	<code>dsmautomig</code>
Start threshold migration for all storage pools in the <code>/home</code> file system.	<code>dsmautomig /home</code>
Start threshold migration for all storage pools in the <code>/home</code> and <code>/test1</code> file systems.	<code>dsmautomig /home /test1</code>
Start threshold migration for the storage pools 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 /fs2 file systems and for the storage pool named gold in the /fs1 file systems. `dsmautomig gold::/fs1 /fs2`

## dsmdf

The **dsmdf** command displays the information for one or more file systems, such as file system state, inode information, and space information.

Specifically, the **dsmdf** command displays information about:

- File system state: active (a), inactive (i), or global inactive (gi)
- Sum of the sizes of all migrated files
- Amount of space that is used on your local file system for premigrated files
- Number of inodes that are used for migrated or premigrated files
- Number of unused inodes on your local file system
- Amount of free space on your local file system

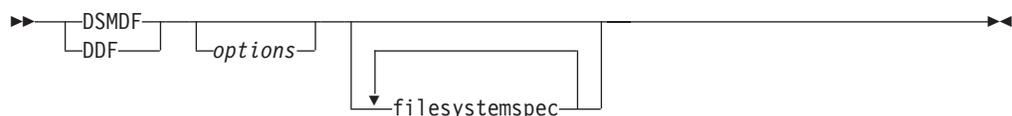
Only migration and recall processes dynamically update status information for your file systems. If any other process changes the state of a file, the change is not reflected in the information that the **dsmdf** command displays until reconciliation is run.

### For AIX GPFS and Linux x86\_64 GPFS file systems only:

Migrating newly created files less than five (5) minutes old might display incorrect results (resident size) when you use the **dsmdf** and **dsmdu** commands. The reason is because GPFS is not synchronized on all nodes when you migrate files. The last block of a file is not released from the disk although the file migrated successfully. This procedure can cause a deviation from an assumed disk usage if many small files are migrated and the block size is high.

**Note:** You can display only information about mounted file systems. If a file system is space managed but not mounted, it does not appear within the command output.

## Syntax



## Parameters

*options*

### **-Detail or when optionformat=short: -v**

Displays information about file systems with each value appearing on its own line. Values representing the amount of space is shown in kilobytes only.

### **-ERRORLOGName=file\_path**

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified

by the **errorlogname** option in options file `dsm.sys` or `dsm.opt`, or as specified by the **DSM\_LOG** environment variable. If any part of the path you specify does not exist, the HSM client creates it.

**-Help or when optionformat=short: -h**

Displays syntax and options for the command. Do not specify other options when you specify the **help** option.

**-Logname=file\_path or when optionformat=short: -Lfile\_path**

Specifies the path and file name of the log file to be used by this command. The kind of events that are logged to the file are specified with the **hsmlogeventflags** option. Error events are not logged to the file.

*filesystemspec*

The name of the file system for which you want to display information. The default is all file systems to which you added space management. You can specify more than one file system name, and you can use wildcard characters within a file system name. If you specify more than one file system name, separate each name with one or more blank spaces.

## Examples

Task	Command
Display information for all file systems to which you added space management.	<code>dsmdf</code>
Display information for the <code>/home</code> file system.	<code>dsmdf /home</code>
Specify log file to be used by the <b>dsmdf</b> command	<code>dsmdf -Logname=mylogfile</code> <code>dsmdf -detail -Logname=/tmp/dsmdflog /home</code> <code>dsmdf -v -Lmylogfile</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_64` GPFS file systems only:

The **dsmdf** and **dsmdu** commands can display incorrect resident file size when you migrate newly-created files. This is because GPFS is not synchronized on all nodes when you migrate files. The last block of a file is not released from the disk although the file migrated successfully. This can cause deviation from assumed disk usage if many small files are migrated and blocksize is high.

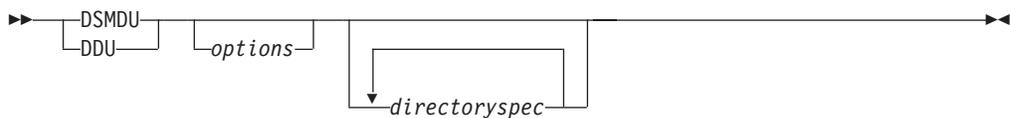
**Attention:** Running **dsm1s** or **dsmdu** on remotely mounted GPFS file systems might show incorrect values. See Table 27 for an example of this situation.

Table 27. Space 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
Space-managed by A	Not Space-managed by B

Only file systems of the local cluster can be managed by HSM. The local Cluster A file system is space-managed by Cluster A, but cannot be space-managed by Cluster B, although it is remotely mounted on Cluster B. The Data Management Application Programming Interface (DMAPI) interface does not work for remotely mounted file systems, so incorrect values might be reported by **dsm1s** 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 when optionformat=short: -a**

Displays the number of 1 KB blocks that each file in the specified directory and each of its subdirectories use.

#### **-ERRORLOGName=file\_path**

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file **dsm.sys** or **dsm.opt**, or as specified by the **DSM\_LOG** environment variable. If any part of the path you specify does not exist, the HSM client creates it.

#### **-Help or when optionformat=short: -h**

Displays syntax and options for the command. Do not specify other options when you specify the **help** option.

#### **-Summary or when optionformat=short: -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>
Display space usage information for each file in the <code>/migfs2/test</code> directory and in all of its subdirectories.	<code>dsmdu -a /migfs2/test</code>
Display the total number of 1 KB blocks that the <code>/migfs2/test</code> directory and all of its subdirectories use.	<code>dsmdu -Summary /migfs2/test</code>

## dsmls

The **dsmls** command displays file information, such as sizes and state.

Specifically, the **dsmls** command displays the following information about a list of files:

- Actual size (in bytes)
- Resident size (in bytes)
- Resident block size (in KB)
- File state and recall mode
- Name of the coupled Tivoli Storage Manager server
- File name

For a resident or premigrated file, the actual size and resident size are the same. For a migrated file, the actual size is the size of the original file. The resident size is the size of the stub file that remains on your local file system.

The file state for a file can 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 notations display in the File State column.

**Attention:**

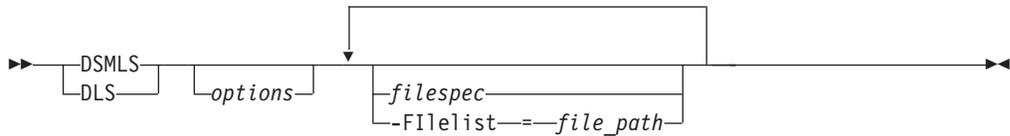
- Running **dsmls** or **dsmdu** on remotely mounted GPFS file systems might show incorrect values. See Table 28 for an example of this situation.

Table 28. Space 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
Space-managed by A	Not Space-managed by B

Only file systems of the local cluster can be managed by HSM. The local Cluster A file system is space-managed by Cluster A, but cannot be space-managed by Cluster B, although it is remotely mounted on Cluster B. The Data Management Application Programming Interface (DMAPI) interface does not work for remotely mounted file systems, so incorrect values might be reported by **dsmls** or **dsmdu**.

**Syntax**



**Parameters**

*options*

**-ERRORLOGName=file\_path**

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file **dsm.sys** or **dsm.opt**, or as specified by the **DSM\_LOG** environment variable. If any part of the path you specify does not exist, the HSM client creates it.

**-Help or when optionformat=short: -h**

Displays syntax and options for the command. Do not specify other options when you specify the **help** option.

**-Noheader or when optionformat=short: -n**

Omits column headings from the output for this command.

**-Recursive or when optionformat=short: -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=file\_path** or when **optionformat=short**: **-fi=file\_path**

Specifies the path of a file that contains a list of files to be processed by this command.

The entries in the list file must adhere to the following rules:

- Each entry is a fully qualified path to a file or a relative path to a file.
- No entry is a directory object.
- Each entry is on a separate line.

The HSM client ignores any entry that does not adhere to these rules.

The following rules apply to a file list for the HSM client. These rules differ from the rules for backup-archive client file lists.

- An entry can contain unprintable characters, but cannot contain a carriage return.
- Wildcard characters are allowed.
- A file name that contains spaces does not have to be enclosed in quotation marks.
- An entry that begins with a quotation mark *and* ends with a quotation mark is tolerated. The HSM client assumes that the quotation marks are not needed and omits them when processing the entry. If beginning and ending quotation marks are required to identify the file, add double quotation marks before the beginning and after the end of the entry.

The following is an example of a list of files within a list file:

```
/home/dir/file1
"/fs1/dir2/file3"
"/fs2/my files/file4"
/fs2/my files/file5
../tivoli/'file1'
"'fs3'/dir3/'file.txt'"
fs4/dir/a"file".txt
'/fs4/dir/file.txt'
/fs5/dir/file*with?wildcards.txt
```

If the list file name that you specify with the **filelist** option does not exist, the command fails.

## Examples

Task	Command
List all files in the current directory.	<code>dsmls</code>
List all files in the <code>/migfs2/test</code> directory.	<code>dsmls /migfs2/test</code>
List all files in the <code>/migfs2/test</code> directory and in its subdirectories.	<code>dsmls -Recursive /migfs2/test</code>
List all files whose names begin with <b>tf</b> in the <code>/migfs2/test</code> directory.	<code>dsmls /migfs2/test/tf*</code>
List all files in the list file named <code>/tmp/filelist</code> .	<code>dsmls -filel=/tmp/filelist</code>

---

## dsmmigfs add and update

Use the **dsmmigfs** command with the **add** parameter to add space management to the file system. Use the **dsmmigfs** command with the **update** parameter to update space management settings for your file system.

You must have root user authority to use this command.

For information about configuring GPFS integration with the HSM client, see the Tivoli field guide *TSM for Space Management for UNIX-GPFS Integration Part I: Policy-driven Threshold Migration* at <https://www.ibm.com/support/docview.wss?uid=swg27018848>.

**For AIX GPFS and Linux x86\_64 GPFS file systems only:** Before you run **dsmmigfs add filesystemName**, ensure that file system is mounted and enabled for Data Management Application Programming Interface (DMAPI) management. Issue the following commands:

For AIX GPFS and Linux x86\_64 GPFS:

```
/usr/lpp/mmfs/bin/mmlsfs DevicePath -z
```

If it is required, change the value to:

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

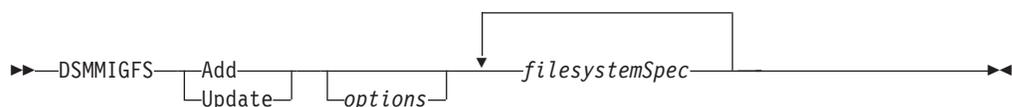
Run only one **dsmmigfs** command within the local GPFS node set at the same time.

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 or when optionformat=short: a**

Adds space management to your file systems.

**Update or when optionformat=short: u**

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.

On GPFS file systems, you can use the GPFS policy engine to monitor space thresholds and search for migration candidates. If you configure the GPFS policy engine to drive automatic migration, the following options for the **dsmmigfs** command have no consequence:

- **hthreshold**
- **lthreshold**
- **maxcandidates**
- **maxfiles**
- **minpartialrecallsize**
- **minmigfilesize**

#### **-Hthreshold=*n* or when optionformat=short: -*tn***

The high threshold percentage that you set for space usage on your file systems. Specify a value of 0 -100 percent. The default is 90 percent.

#### **-Lthreshold=*n* or when optionformat=short: -*ln***

The low threshold percentage that you set for space usage on your file systems. Specify a value of 0 through 100 percent. The default is 80 percent.

#### **-ERRORLOGName=*file\_path***

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file **dsm.sys** or **dsm.opt**, or as specified by the **DSM\_LOG** environment variable. If any part of the path you specify does not exist, the HSM client creates it.

#### **-LOGname=*file\_path* or when optionformat=short: -*Lfile\_path***

Specifies the path and file name of the log file to be used by this command. The kind of events that are logged to the file are specified with the **hsmlogeventflags** option. Error events are not logged to the file.

#### **-Maxcandidates=*n* or when optionformat=short: -*mn***

The maximum number of migration candidates the **dsmscoutd** daemon delivers to the automatic migration process during automigration. A value 50 - 1000 is sufficient for good performance. The range of values is 9 through 9999999. The default is 10000.

**-MAXFiles=*n* or when optionformat=short: -*cn***

The maximum number of files for which the CFI database is sized. The range of values is 0 through 4294967295. The default is 0.

This parameter allocates space based on the maximum number of blocks in the file system. If the specified value is less than the number of currently used blocks, an error occurs. However, only these items are honored:

- 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

You can configure the size of the complete file index (CFI). This value can be queried or set with the `maxfiles` parameter. If this parameter is set to 0, the CFI allocates the maximum required space, which is equated to the maximum number of blocks possible in the file system. Less file system space is needed to create the CFI. If the specified value is out of bounds, the nearest boundary is used to allocate the configurable CFI size.

**-MINMigfilesize=*n* or when optionformat=short: -*yn***

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, or the value is ignored. A valid value takes precedence over the global `minmigfilesize` option setting specified in the `dsm.sys` file.

**-MINPartialrecallsize=*n* or when optionformat=short: -*nn* (AIX GPFS and Linux x86\_64 GPFS only)**

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=*n* or when optionformat=short: -*zn***

Specifies a number to enable or disable an asynchronous recall of migrated files. This parameter is the number of megabytes that must be recalled before HSM starts streaming data to the requesting application (to ensure a steady stream of data). Smaller files are recalled using normal recall mode. 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=*n* or when optionformat=short: -*pn***

The percentage of file system space that is available to contain premigrated files. The minimum value is 0. The maximum value is the low threshold. A warning is issued if the premigration percentage is greater than the difference between the low threshold and the recommended minimum size of the file system. A premigration percentage greater than this difference can result in insufficient space for the file system, and can result in endless attempts to find premigration candidates.

**-Quota=*n* or when optionformat=short: -*qn***

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.

**-Server=server\_name or when optionformat=short: -Sservername**

Overrides the default migration server for this file system. Specify the server to contact for space management services. Define the server in a stanza in your dsm.sys file. If you do not specify a server name, the default migration server that you defined is used. Use a dash (-) to set the server to the default migration server.

**-Stubsize=n or when optionformat=short: -sn**

The size of stub files remaining on your local file systems when files migrate to storage. The following values can be assigned:

0|(16|32|64|128|256|512)(K|M)|1G

For HSM clients on AIX GPFS and Linux x86\_64 GPFS file systems, you can specify 0 or a multiple of the file system fragment size. The default value is 0. A stub file on a GPFS file system occupies at least 8 KB because of the metadata in the stub file.

If HSM is globally deactivated on a node (**dsmmigfs globaldeactivate**), it is reset to active state, if one of the following commands are performed:

```
dsmmigfs add
dsmmigfs remove
dsmmigfs update
dsmmigfs takeover
dsmmigfs rollback
dsmmigfs globalreactivate
```

## Examples

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 /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>
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"> <li>• Change the high threshold to 80 percent.</li> <li>• Change the low threshold to 70 percent.</li> <li>• Set the size of stub files to 1 megabyte.</li> </ul>	<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>
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>

## Related tasks

“Updating settings from the HSM GUI” on page 47

## Related reference

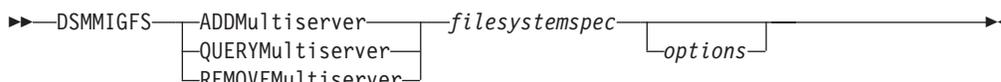
“minmigfilesize” on page 111

---

## dsmmigfs addmultiserver, querymultiserver, and removemultiserver

Use the **dsmmigfs** command with the **addmultiserver**, **querymultiserver**, or **removemultiserver** parameter to manage space in an environment with multiple Tivoli Storage Manager servers.

### Syntax



### Parameters

#### **ADDMultiserver** or when **optionformat=short: am**

Adds a Tivoli Storage Manager server to the list of servers that can manage the file system.

#### **QUERYMultiserver** or when **optionformat=short: qm**

Queries the status of the multiple-server environment. The command prints all Tivoli Storage Manager servers that can manage the specified file system. For each server, the output displays the number of files, number of bytes, and migration throughput.

#### **REMOVEMultiserver** or when **optionformat=short: rm**

Removes an existing Tivoli Storage Manager server from the list of servers that can manage the specified file system. That Tivoli Storage Manager server is no longer a target for migration and backup. Current<sup>®</sup> migration and backup processes are not affected. Removal does not affect the recall or restore of files migrated or backed up to the Tivoli Storage Manager server.

After removing the Tivoli Storage Manager server from the list of servers, and before physically removing the server, run the **dsmRemoveServer.pl** script. The **dsmRemoveServer.pl** script recalls all files from the Tivoli Storage Manager server to the local file systems. After running the **dsmRemoveServer.pl** script, you can migrate and back up files to another Tivoli Storage Manager server.

#### *filesystemspec*

The file system name on which to perform the specified action. Specify only one file system.

#### *options*

##### **-ERRORLOGName=***file\_path*

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file **dsm.sys** or **dsm.opt**, or as specified by the **DSM\_LOG** environment variable. If any part of the path you specify does not exist, the HSM client creates it.

##### **-Logname=***file\_path* or when **optionformat=short: -L***file\_path*

Specifies the path and file name of the log file to be used by this

command. The kind of events that are logged to the file are specified with the **hsmlogeventflags** option. Error events are not logged to the file.

**-Server=server\_name or when optionformat=short: -sserver\_name**

Specify the target server for the task. This option is valid only in a multiple-server environment. If the file system is managed by multiple servers and you do not specify **server**, the task is attempted with the default migration server.

If a file is coupled with a server, the value of this option must be that coupled server. If you specify another Tivoli Storage Manager server, the task fails.

#### Related tasks

“Enabling a file system to be managed by multiple Tivoli Storage Manager servers” on page 54

“Removing a Tivoli Storage Manager server from a multiple-server environment” on page 56

#### Related reference

“dsmMultiServerUpgrade.pl” on page 167

“dsmRemoveServer.pl” on page 169

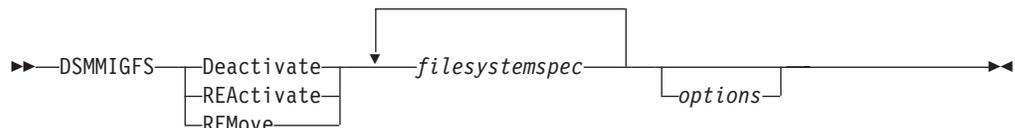
---

## 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 or when optionformat=short: -dea**

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 or when optionformat=short: -rea**

Reactivates space management for a file system.

#### **REMove or when optionformat=short: -rem**

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



command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file `dsm.sys` or `dsm.opt`, or as specified by the **DSM\_LOG** environment variable. If any part of the path you specify does not exist, the HSM client creates it.

**-Logname=*file\_path* or when optionformat=short: -L*file\_path***  
 Specifies the path and file name of the log file to be used by this command. The kind of events that are logged to the file are specified with the **hsmlogeventflags** option. Error events are not logged to the file.

## Examples

Task	Command
Globally deactivate space management for your client node.	<code>dsmmigfs GLOBALDeactivate</code>
Globally reactivate space management for your client node.	<code>dsmmigfs GLOBALReactivate</code>

## dsmmigfs help

Use the **dsmmigfs** command with the **help** parameter to display the **dsmmigfs** command syntax and options.

The syntax shown is appropriate for the value of **optionformat**. For example, when **optionformat=short**, the short-format syntax is shown.

### Syntax

►► `DSMMIGFS—Help` options ◀◀

### Parameters

#### Help or when optionformat=short: h

Displays the syntax of the **dsmmigfs** command including command parameters, options, and valid ranges of options values.

#### options

##### **-ERRORLOGName=*file\_path***

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file `dsm.sys` or `dsm.opt`, or as specified by the **DSM\_LOG** environment variable. If any part of the path you specify does not exist, the HSM client creates it.

##### **-Logname=*file\_path* or when optionformat=short: -L*file\_path***

Specifies the path and file name of the log file to be used by this command. The kind of events that are logged to the file are specified with the **hsmlogeventflags** option. Error events are not logged to the file.

## Examples

Task	Command
Display the syntax of the <b>dsmmigfs</b> command.	<code>dsmmigfs h</code>

## dsmmigfs query

Use the **dsmmigfs** command used with the **query** parameter to display the current space management settings for a file system.

### Syntax



### Parameters

#### Query

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
- Maximum files
- Server name

**Note:** These settings are displayed only if you specify the **detail** option.

#### *options*

##### **-Detail or when optionformat=short: -v**

Displays HSM settings for each file system. Displays additional 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. )

- Maximum candidates
- Partial file recall size
- Streaming recall size
- The minimum size (in bytes) that a file must qualify for migration

For AIX GPFS and Linux 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.**

This is the node name for each node ID and the frame ID for both the owner and the preferred nodes.

**Preferred node.**

This is the cluster node where HSM was initially added to the file system or the node where the HSM administrator has performed a **dsmmigfs takeover** operation.

**Owner node.**

This is the cluster node that is currently managing the file system (it can be different from the preferred node after failover).

The normal `dsmmigfs query` command does not provide any GPFS-specific information (no node set ID).

**Tip:** The `dsmmigfs query -detail` command displays the status of the local failover environment (either active or inactive).

**-Failover or when optionformat=short: -f**

Provides a status overview of the failover environment of all HSM managed cluster nodes. This choice is valid for AIX GPFS and Linux 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 applying 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.

**-ERRORLOGName=file\_path**

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file `dsm.sys` or `dsm.opt`, or as specified by the **DSM\_LOG** environment variable. If any part of the path you specify does not exist, the HSM client creates it.

**-Logname=file\_path or when optionformat=short: -Lfile\_path**

Specifies the path and file name of the log file to be used by this command. The kind of events that are logged to the file are specified with the **hsmlogeventflags** option. Error events are not logged to the file.

*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 space management settings for the /migfs2 file system.	<code>dsmmigfs query /migfs2</code>
Display the space management entries made in the log file /tmp/migfslog.	<code>dsmmigfs query -L/tmp/migfslog /migfs2</code>
Display the space management entries made in the log file /tmp/migfslog.	<code>dsmmigfs query -Logname=/tmp/migfslog /migfs2</code>

## Related concepts

Chapter 3, “Configuring the HSM client,” on page 23

## Related reference

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

---

## dsmmigfs rollback

Use the **dsmmigfs** command with the **rollback** parameter to transfer the space management of file systems to the preferred node if the node is different from the current owner node.

Valid on AIX GPFS and Linux x86\_64 GPFS file systems only.

You must have root user authority to use this command.

## Syntax

```
▶▶—DSMMIGFS—ROLLback—options—▶▶
```

## Parameters

### ROLLback or when optionformat=short: **-roll**

Transfers the HSM management of file systems to the preferred node if the node is different from the current owner node. Enter this command on the preferred node.

### *options*

#### **-ERRORLOGName=***file\_path*

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file *dsm.sys* or *dsm.opt*, or as specified by the **DSM\_LOG** environment variable. If any part of the path you specify does not exist, the HSM client creates it.

#### **-Logname=***file\_path* or when optionformat=short: **-L***file\_path*

Specifies the path and file name of the log file to be used by this command. The kind of events that are logged to the file are specified with the **hsmlogeventflags** option. Error events are not logged to the file.

## Examples

Task	Command
Transfer to the preferred node.	<code>dsmmigfs rollback</code>

---

## dsmmigfs sdrreset, enablefailover, and disablefailover

Use the **dsmmigfs** command used with the **sdrreset**, **enablefailover**, or **disablefailover** parameter to manage recovery from partial system failure.

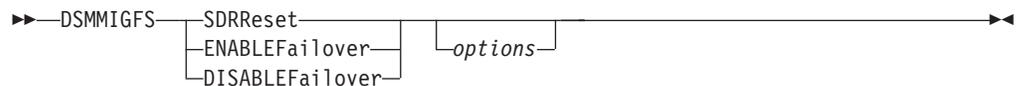
Valid on AIX GPFS and Linux x86\_64 GPFS file systems only.

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

#### **SDRReset** or when **optionformat=short: SDRr**

**Do not use this command during normal operation.**

The HSM client resets 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** or when **optionformat=short: enablef**

Activates the node for failover operations within the GPFS cluster.

#### **DISABLEFailover** or when **optionformat=short: disablef**

Deactivates failover operations on the node.

#### *options*

##### **-ERRORLOGName=***file\_path*

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file `dsm.sys` or `dsm.opt`, or as specified by the **DSM\_LOG** environment variable. If any part of the path you specify does not exist, the HSM client creates it.

##### **-Logname=***file\_path* or when **optionformat=short: -L***file\_path*

Specifies the path and file name of the log file to be used by this command. The kind of events that are logged to the file are specified with the **hsmlogeventflags** option. Error events are not logged to the file.

## Examples

Task	Command
Deactivate failover operations on the HSM client node.	<code>dsmmigfs disableFailover</code>

## dsmmigfs stop, start, and restart

use the **dsmmigfs** command with the **stop**, **start**, or **restart** parameter to control space management daemons.

Valid on AIX GPFS, Linux GPFS systems only

You must have root user authority to use this command.

Use the **dsmmigfs** command with the **stop**, **start**, or **restart** parameter to control daemons:

- Start all daemons
- Stop all daemons, **dsmrecall** and **dsmmigrate** processes except **dsmwatchd**
- Restart all daemons, stop **dsmrecall** and **dsmmigrate** except **dsmwatchd**.

**Note:** Be aware that the daemons will be started with the same environment as the **dsmwatchd** daemon, which means that options files `dsm.opt` and `dsm.sys` in the default installation path `/usr/tivoli/tsm/client/ba/bin` will be used.

### Syntax



### Parameters

#### **START** or when **optionformat=short: start**

Starts all HSM daemons on the local client node. The **dsmwatchd** daemon is not affected.

#### **STOP** or when **optionformat=short: stop**

Stops all HSM daemons. The **dsmrecall** and **dsmmigrate** processes are stopped. The **dsmwatchd** daemon is not affected.

#### **RESTART** or when **optionformat=short: restart**

Restarts all HSM daemons. The **dsmrecall** and **dsmmigrate** processes are started. The **dsmwatchd** daemon is not affected.

#### *options*

##### **-ERRORLOGName=***file\_path*

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file `dsm.sys` or `dsm.opt`, or as specified by the **DSM\_LOG** environment variable. If any part of the path you specify does not exist, the HSM client creates it.

##### **-Logname=***file\_path* or when **optionformat=short: -L***file\_path*

Specifies the path and file name of the log file to be used by this

command. The kind of events that are logged to the file are specified with the **hsmlogeventflags** option. Error events are not logged to the file.

## Examples

Task	Command
Start all daemons	<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> options files.	<code>dsmmigfs RESTART</code>

---

## dsmmigfs takeover

Use the **dsmmigfs** command with the **takeover** parameter to transfer the HSM management of a file system to another the HSM client node within the same local GPFS node set.

Valid on AIX GPFS and Linux x86\_64 GPFS file systems only.

You must have root user authority to use this command.

## Syntax

►► `DSMMIGFS TAKEover filespec` options ◀◀

## Parameters

### TAKEover or when optionformat=short: take

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.

### options

#### **-ERRORLOGName=file\_path**

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file `dsm.sys` or `dsm.opt`, or as specified by the **DSM\_LOG** environment variable. If any part of the path you specify does not exist, the HSM client creates it.

#### **-Logname=file\_path** or when optionformat=short: **-Lfile\_path**

Specifies the path and file name of the log file to be used by this command. The kind of events that are logged to the file are specified with the **hsmlogeventflags** option. Error events are not logged to the file.

## Examples

Task	Command
Transfer the HSM management of the current directory to the 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—[options]—▶▶
```

### Parameters

*options*

**-ERRORLOGName=***file\_path*

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file `dsm.sys` or `dsm.opt`, or as specified by the **DSM\_LOG** environment variable. If any part of the path you specify does not exist, the HSM client creates it.

### 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.



Display information about management classes that you can assign to files on your client node. `dsmmigquery -Mgmtclass -Detail`

#### Related tasks

“Updating settings from the HSM GUI” on page 47

## 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.

See *IBM Tivoli Storage Manager for UNIX and Linux Backup-Archive Clients Installation and User's Guide* for information about using these shared options.

- `asnodename`
- `commethod`
- `compression`
- `defaultserver`
- `detail`
- `diskbuffsize`
- `enablelanfree`
- `errorlogmax`
- `errorlogname`
- `errorlogretention`
- `exclude`
- `exclude.compression`
- `incl excl`
- `include`
- `include.compression`
- `lanfreecommmethod`
- `lanfreetcpport`
- `lanfreeshmport`
- `makesparsefile`
- `nodename`
- `passwordaccess`
- `passworddir`
- `servername`
- `shmport`
- `skipacl`
- `tcpbuffsize`
- `tcpnodelay`
- `tcpport`
- `tcpserveraddress`
- `tcpwindowsize`

---

## dsmmigrate

The **dsmmigrate** command selects specific files from the local file system and migrates them to a Tivoli Storage Manager server.

**Attention:** On large file systems, selective migration can take much time.

If the file exists in a file system with a different server stanza than the last processed file, a new session starts for each migrated file. This action can happen as a result of links from one file system to another.

### Note:

1. If you set the **tapeprompt** option to YES in the `dsm.opt` file, and the destination for the 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 HSM client does not migrate contents of symbolic links. Symbolic links are not followed during recursive selective migration.

The first file migrates even if the file size exceeds the quota that you specified for the file system. When migration occurs, the **ddf** command displays zero migrated and premigrated bytes for the file system. If the total number of bytes exceeds the quota after the file migrates, the next file is not migrated.

## Syntax



## Parameters

### options

#### **-Detail or when optionformat=short: -v**

Displays the size and file name for each file that you migrate.

#### **-ERRORLOGName=file\_path**

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file `dsm.sys` or `dsm.opt`, or as specified by the **DSM\_LOG** environment variable. If any part of the path you specify does not exist, the HSM client creates it.

#### **-Help or when optionformat=short: -h**

Displays syntax and options for the command. Do not specify other options when you specify the **help** option.

#### **-Logname=file\_path or when optionformat=short: -Lfile\_path**

Specifies the path and file name of the log file to be used by this command. The kind of events that are logged to the file are specified with the **hsmlogeventflags** option. Error events are not logged to the file.

#### **-Premigrate or when optionformat=short: -p**

Files that are migrated are in the premigrated state.

**-Recursive or when optionformat=short: -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.

The HSM client does not migrate contents of symbolic links. Symbolic links are not followed during recursive selective migration.

**-Server=server\_name or when optionformat=short: -server\_name**

Specify the target server for the task. This option is valid only in a multiple-server environment. If the file system is managed by multiple servers and you do not specify **server**, the task is attempted with the default migration server.

If a file is coupled with a Tivoli Storage Manager server, you must specify that server. If a file has not been coupled with a Tivoli Storage Manager server, you can specify any Tivoli Storage Manager server that was added to the multiple-server environment.

*filespec*

The path and file name of the files that you want to migrate. This parameter is required. You can use wildcard characters to specify a group of files or all files in a directory. You can enter more than one file specification in one command. If you enter more than one file specification, separate each specification with one or more blank spaces.

**-Filelist=file\_path or when optionformat=short: -fi=file\_path**

Specifies the path of a file that contains a list of files to be processed by this command.

The entries in the list file must adhere to the following rules:

- Each entry is a fully qualified path to a file or a relative path to a file.
- No entry is a directory object.
- Each entry is on a separate line.

The HSM client ignores any entry that does not adhere to these rules.

The following rules apply to a file list for the HSM client. These rules differ from the rules for backup-archive client file lists.

- An entry can contain unprintable characters, but cannot contain a carriage return.
- Wildcard characters are allowed.
- A file name that contains spaces does not have to be enclosed in quotation marks.
- An entry that begins with a quotation mark *and* ends with a quotation mark is tolerated. The HSM client assumes that the quotation marks are not needed and omits them when processing the entry. If beginning and ending quotation marks are required to identify the file, add double quotation marks before the beginning and after the end of the entry.

The following is an example of a list of files within a list file:

```
/home/dir/file1
"/fs1/dir2/file3"
"/fs2/my files/file4"
/fs2/my files/file5
../tivoli/'file1'
```

```
|
| "'fs3'/dir3/'file.txt'"
| fs4/dir/a"file".txt
| '/fs4/dir/file.txt'
| /fs5/dir/file*with?wildcards.txt
```

If the list file name that you specify with the **filelist** option does not exist, the command fails.

## Examples

Task	Command
Migrate all files in a directory 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. Log results in the file /tmp/miglog.	<code>dsmmigrate -Detail -L/tmp/miglog tf04</code>
Migrate all files in a filelist named /tmp/filelist. Log results in the file /tmp/miglog.	<code>dsmmigrate -logname=/tmp/miglog -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 a stub file or an original copy of a premigrated file is deleted from your local file system, the corresponding migrated or premigrated file is marked for expiration when reconciliation runs again.

**Note:** The **dsmmigundelete** command does not support hard-linked files. If you recreate a stub file for a hard-linked file, a stub file is not recreated *unless* all of the hard-linked files are deleted from your local file system. When one file in a set of hard-linked files is migrated, all of the hard-linked files in the set become stub files. When the **dsmmigundelete** command recreates a stub file for a hard-linked file, the stub file has the same name as the file that was originally migrated. Stub files are not recreated for any other files that were previously in the hard-linked set of files.

## Syntax

```
▶▶—DSMMIGUNDELETE—┬───┬───▶
 └───┬───┬───┘
 [filesystemspec] [options]
```

## Parameters

*options*

**-Recover or when optionformat=short: -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 overwrites all premigrated or recalled files. Data is lost from the last

migration up to when this option is used! After the **dsmmigundelete -recover** command is issued, there is no way you can check whether the data of the local stub file is different from the contents of the file in the Tivoli Storage Manager server.

Use this option only if all or most of the files in the file system are stub files and you need good performance during stub file recreation.

For performance purposes, do not use this option with the **detail** option.

Files that were never migrated are not recreated.

**-Detail or when optionformat=short: -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 when optionformat=short: -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.

**-ERRORLOGName=file\_path**

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file **dsm.sys** or **dsm.opt**, or as specified by the **DSM\_LOG** environment variable. If any part of the path you specify does not exist, the HSM client creates it.

**-Logname=file\_path or when optionformat=short: -Lfile\_path**

Specifies the path and file name of the log file to be used by this command. The kind of events that are logged to the file are specified with the **hsmlogeventflags** option. Error events are not logged to the file.

**-Server=server\_name or when optionformat=short: -sserver\_name**

Specify the target server for the task. This option is valid only in a multiple-server environment. If the file system is managed by multiple servers and you do not specify **server**, the task is attempted with the default migration server.

If a file is coupled with a server, the value of this option must be that coupled server. If you specify another Tivoli Storage Manager server, the task fails.

### *filesystemspec*

The name of the file system for which you want to recreate deleted stub files and create stub files for premigrated files that were deleted from your local file system. The default is all file systems for which space management is active. You can specify more than one file system name. If you specify several file system names, separate each name with one or more blank spaces.

## Examples

Task	Command
Recreate stub files for migrated files 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. Results are logged in the log file /tmp/undeletelog.	dsmmigundelete -expiring -Logname=/tmp/undeletelog /home
Recreate stub files for migrated files in the /home file system not   marked for expiration. Reconciliation was not run because the files   were deleted. Entries are made in the log file /tmp/undeletelog.	dsmmigundelete -L/tmp/undeletelog /home
Recreates stub files for migrated files that had been removed from the /trullofs file system.	dsmmigundelete -recover /trullofs

## dsmonitord

The **dsmonitord** command starts the HSM space monitor daemon if it has stopped. If you issue this command and the space monitor daemon is running, action is not taken.

You must have root user authority to use this command.

### Syntax

▶▶—DSMMONITORD—options▶▶

### Parameters

#### *options*

#### **-ERRORLOGName=***file\_path*

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file dsm.sys or dsm.opt, or as specified by the **DSM\_LOG** environment variable. If any part of the path you specify does not exist, the HSM client creates it.

#### **-Help or when optionformat=short: -h**

Displays syntax and options for the command. Do not specify other options when you specify the **help** option.

---

## 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 coupled Tivoli Storage Manager server
- 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—┐  
          └options┘

### Parameters

*options*

**-ERRORLOGName=***file\_path*

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file `dsm.sys` or `dsm.opt`, or as specified by the **DSM\_LOG** environment variable. If any part of the path you specify does not exist, the HSM client creates it.

**-Help or when optionformat=short: -h**

Displays syntax and options for the command. Do not specify other options when you specify the **help** option.

### Examples

Task	Command
Display the status of recall processes.	<code>dsmq</code>

## dsmrecall

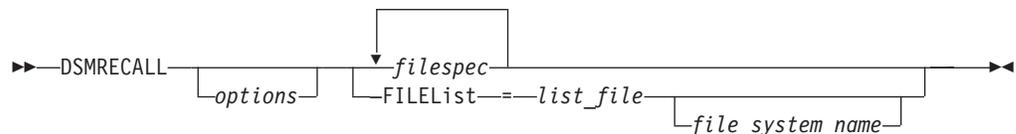
The **dsmrecall** command selectively recalls migrated files or parts of migrated files to the local file system. Space management must be active.

**Note:** On large file systems, selective recall can take a long time.

If the file resides in a file system whose server stanza is different from the stanza of the last processed file, a new session is started for each recalled file. This process can happen as a result of links from one file system to another.

To display a list of all migrated files, use the **dsmmigquery** command. To display information about a list of migrated files from a particular file system or directory, use the **dsm ls** command.

### Syntax



### Parameters

#### *options*

##### **-Detail or when optionformat=short: -v**

Display the size, path, and file name for each file that you recall.

##### **-ERRORLOGName=file\_path**

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file `dsm.sys` or `dsm.opt`, or as specified by the **DSM\_LOG** environment variable. If any part of the path you specify does not exist, the HSM client creates it.

##### **-Help or when optionformat=short: -h**

Displays syntax and options for the command. Do not specify other options when you specify the **help** option.

##### **-OFFSET=n or when optionformat=short: -on**

Specify the offset from the beginning of the file of the required data range for partial recall (in bytes, kilobytes, megabytes, or gigabytes). Multipliers (k, m, g, K, M, or G) can be used. The range of values is 0 through 2147483647. There is no default value. This option is valid on AIXGPFS and Linux x86\_64 GPFS file systems only.

**Remember:** Use this option only with the **size** option and when a file system is specified. The **recursive**, **detail**, and **filelist** options are not valid with the **offset** option.

Examples are: `-offset=10` (bytes), `-offset=23k` (kilobytes), `-o5M` (megabytes), `-o2G` (gigabytes).

##### **-PREview or when optionformat=short: -p**

Generate list files that are optimized for tape recalls but do not recall the

files. You must also specify **filelist** and a file system. The **preview** option is not valid when **filelist** specifies a collection file.

**-Logname=*file\_path* or when optionformat=short: -L*file\_path***

Specifies the path and file name of the log file to be used by this command. The kind of events that are logged to the file are specified with the **hsmlogeventflags** option. Error events are not logged to the file.

**-Recursive or when optionformat=short: -R**

Recall files that match the file specification in a directory and its subdirectories. If you do not use this option, files are recalled only for those directories that you specify.

**-SIZE=*n* or when optionformat=short: -s*n***

Specify the size of the required data range for partial recall (in bytes, kilobytes, megabytes, or gigabytes). Multipliers (k, m, g, K, M, or G) can be used. The range of values is 0 - 4294967295. There is no default value. This option is valid on AIXGPFS and Linux 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).

**-Server=*server\_name***

Specify the target server for the task. This option is valid only in a multiple-server environment. If the file system is managed by multiple servers and you do not specify **server**, the task is attempted with the default migration server.

If a file is coupled with a server, the value of **server** must be that coupled server. If you specify another Tivoli Storage Manager server, the task fails with a warning message.

### *filespec*

Specify the path and file name of the files to recall. This parameter is required. You can use wildcard characters to specify a group of files or all files in a directory, or more than one file specification in one command. When you use wildcard characters in a file specification, the HSM client recalls all files that match the specification. If a file matches the specification but it is not migrated, an error message displays. If you enter more than one file specification, separate each specification with at least one blank space.

**-FILEList=*list\_file***

Process the files that are listed in *list\_file*.

The entries in the list file must adhere to the following rules:

- Each entry is a fully qualified path to a file or a relative path to a file.
- No entry is a directory object.
- Each entry is on a separate line.

The HSM client ignores any entry that does not adhere to these rules.

The following is an example of a list of files within a list file:

```
/home/dir/file1
"/fs1/dir2/file3"
"/fs2/my files/file4"
/fs2/my files/file5
../tivoli/'file1'
```

```
|
| "'fs3'/dir3/'file.txt'"
| fs4/dir/a"file".txt
| '/fs4/dir/file.txt'
| /fs5/dir/file*with?wildcards.txt
|
```

If the list file name that you specify with the **filelist** option does not exist, the command fails.

You can specify a collection file generated by **dsmrecall** or a list file generated by another application. You can specify only one list file or collection file, and you can specify the **filelist** option only once in each command.

**If you also specify the preview option:**

You must also specify *file\_system\_name*. The HSM client generates tape-optimized list files and a collection file. No files are recalled.

**If you do not specify the preview option:**

The HSM client recalls files.

**If you also specify *file\_system\_name*:**

The HSM client performs a tape-optimized recall.

**If the list file is a collection file that was generated by **dsmrecall**:**

Recall starts. The list files referenced by the collection are processed in the order defined in the collection file. You must specify the collection file generated by **dsmrecall**; you cannot specify a tape list file or disk list file.

**If the list file was not generated by **dsmrecall**:**

The HSM client first reorders the list to optimize recalls from tape. Then the HSM client recalls the files.

**If you do not specify *file\_system\_name*:**

The HSM client recalls the files in the list file. The HSM client does not reorder the list to optimize recalls from tape.

*file\_system\_name*

Specify a file system. You must also specify the **filelist** option. The HSM client optimizes the processing of files stored on tape.

## Examples

Task	Command
Recall a single file named /migfs1/test/tf04 and display detailed information.	dsmrecall -Detail /migfs1/test/tf04
Recall all migrated files in a directory named /mfs4/user1 and all migrated files in its subdirectories.	dsmrecall -Recursive /mfs4/user1/*
Recall all files in the FILElist named /tmp/filelist. Results are logged in the log file /tmp/recalllog.	dsmrecall -Logname=/tmp/recalllog -filelist=/tmp/filelist
Recall the specified portion of a file named /mfs1/file10 in partial recall mode.	dsmrecall -offset=10M -size=500M /mfs1/file10

### Related concepts

“List files for optimized tape recalls” on page 81

### Related reference

“dsmls” on page 129

“dsmmigquery” on page 146

---

## dsmrecalld

The **dsmrecalld** command starts a recall daemon if it is not running.

You must have root user authority to use this command.

### Note:

- No action is taken if you issue this command while a recall daemon is running.
- When using the backup-archive client to restore space-managed files, **dsmrecalld** must be running.

## Syntax

```
▶▶—DSMRECALLD—┬── options ─┬──▶▶
```

## Parameters

### *options*

**—ERRORLOGName=***file\_path*

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file *dsm.sys* or *dsm.opt*, or as specified by the **DSM\_LOG** environment variable. If any part of the path you specify does not exist, the HSM client creates it.

**—Help or when optionformat=short: —h**

Displays syntax and options for the command. Do not specify other options when you specify the **help** option.

---

## dsmreconcile

The **dsmreconcile** command synchronizes the file systems on your client node with the 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—  
|  
| options |  
|

## Parameters

*options*

**—Detail or when optionformat=short: -d**

Prints progress messages while receiving the list of migrated files from the server.

**—ERRORLOGName=*file\_path***

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file *dsm.sys* or *dsm.opt*, or as specified by the **DSM\_LOG** environment variable. If any part of the path you specify does not exist, the HSM client creates it.

**—FILEINFO or when optionformat=short: -f**

Performs the basic reconciliation task of expiring and deleting obsolete copies from the server. This option is the default.

**—FILELIST=*file\_path* or when optionformat=short: -fi=*file\_path***

Specifies the path of a file that contains a list of files to be processed by this command.

**—Help or when optionformat=short: -h**

Displays syntax and options for the command. Do not specify other options when you specify the **help** option.

**—ORPHANCHECK or when optionformat=short: -o**

Checks for orphaned files in the local file system.

The HSM client verifies that valid objects for each local stub file exist in the space management pool of the Tivoli Storage Manager server. The **dsmreconcile** process determines all migrated and premigrated files, and checks that corresponding objects exist on the server. When orphans are located, their names are recorded in the *file\_system\_name/.SpaceMan/orphan.stubs* file. This command option overrides the value of the **checkfororphans** option in the *dsm.sys* file.

If you previously ran the **dsmreconcile** command with the **preptwo** option, Tivoli Storage Manager performs a two-way orphan check.

The two-way orphan check processing identifies orphans on the file system and on the Tivoli Storage Manager server. The two-way check is done in a single pass and the orphan identification process uses parallel processing.

**Tip:** If reconciliation is controlled by the HSM client space monitor daemon, the scout daemon must be running when you specify this option. If reconciliation is controlled by the GPFS policy engine, the scout daemon is not necessary.

When the reconciliation process is running in orphan-check mode, metadata information is updated on the server. Obsolete copies of files are not expired or deleted on the server.

| **-PREPTWO or when optionformat=short: -p**

| This option prepares a file system for a two-way orphan check. Use this  
| option when reconciliation is controlled by the GPFS policy engine.

| When reconciliation is controlled by the GPFS policy engine, migrated and  
| premigrated files are associated with DMAPI attribute **extObjId**. If you set  
| **hsmextobjidattr=YES**, the DMAPI attribute **extObjId** is associated with the  
| files when they are migrated or premigrated. The **preptwo** option makes  
| the association for any files that were not associated with the **extObjId**  
| attribute when the files were migrated or premigrated.

| The **preptwo** option is valid only if **hsmextobjidattr=YES**.

| When you specify the **preptwo** option, do not specify any of the following  
| options:

- | • **fileinfo**
- | • **filelist**
- | • **orphancheck**

| **-Server=server\_name or when optionformat=short: -sserver\_name**

| Specify the target server for the task. This option is valid only in a  
| multiple-server environment. If the file system is managed by multiple  
| servers and you do not specify **server**, the task is attempted with the  
| default migration server.

| If a file is coupled with a server, the value of this option must be that  
| coupled server. If you specify another Tivoli Storage Manager server, the  
| task fails.

| *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.	dsmreconcile
Start reconciliation for the /migfs1 file system.	dsmreconcile /migfs1
Start an orphan check reconciliation process for the /migfs1 file system.	dsmreconcile -o /migfs1
Start reconciliation for file systems /home and /test1.	dsmreconcile /home /test1
Prepare the two-way orphan check reconciliation for the /migfs1 file system.	dsmreconcile -p /migfs1
Start the policy-based reconciliation on the files listed in the /home/dsn/filelist file list.	dsmreconcile -o -filelist=/home/dsn/filelist

### Related concepts

“The scout daemon” on page 90

---

## dsmrm

The **dsmrm** command removes a recall process from the queue. To obtain the required recall ID to remove a recall process, use the **dsmq** command.

You must have root user authority to use this command.

**Note:** After a recall process starts, issue the **dsmrm** command to stop the process. Do not use **Ctrl+C** to stop a recall process.

### Syntax



### Parameters

*recallid*

The recall process ID that you want to remove from the queue.

*options*

**-ERRORLOGName=***file\_path*

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file **dsm.sys** or **dsm.opt**, or as specified by the **DSM\_LOG** environment variable. If any part of the path you specify does not exist, the HSM client creates it.

**-Help or when optionformat=short: -h**

Displays syntax and options for the command. Do not specify other options when you specify the **help** option.

### Examples

Task	Command
Remove recall ID 10 from the queue.	<code>dsmrm 10</code>
Remove recall ID 5 and recall ID 6 from the queue.	<code>dsmrm 5 6</code>

---

## 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**
- **dsmis**
- **dsmmigfs query**
- **dsmmighelp**
- **dsmmigrate**

- `dsmrecall`

**Note:** If you issue this command while a root daemon is running, no action is taken.

## Syntax



## Parameters

*options*

**-ERRORLOGName=***file\_path*

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file `dsm.sys` or `dsm.opt`, or as specified by the **DSM\_LOG** environment variable. If any part of the path you specify does not exist, the HSM client creates it.

**-Help or when optionformat=short: -h**

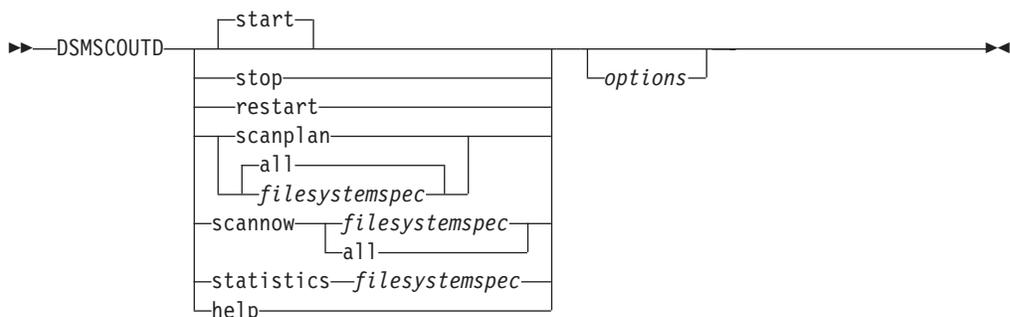
Displays syntax and options for the command. Do not specify other options when you specify the **help** option.

## dsmscoutd

The **dsmscoutd** command starts, stops, and restarts the scout daemon. You can also use the command to start a scan, show future scan times, and show past scan information for file systems.

You must have root user authority to use this command.

## Syntax



## Parameters

All parameters must be entered as indicated. The parameters must be entered in lower case. The parameters have no abbreviation. The parameters have no short format.

**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 space-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.

*options*

| **-ERRORLOGName=***file\_path*

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file `dsm.sys` or `dsm.opt`, or as specified by the **DSM\_LOG** environment variable. If any part of the path you specify does not exist, the HSM client creates it.

| **-Help or when optionformat=short: -h**

Displays syntax and options for the command. Do not specify other options when you specify the **help** option.

---

## 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—┬──options──┬──oldpassword──newpassword──▶
```

## Parameters

### *options*

**-ERRORLOGName=***file\_path*

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file `dsm.sys` or `dsm.opt`, or as specified by the **DSM\_LOG** environment variable. If any part of the path you specify does not exist, the HSM client creates it.

**-Help or when optionformat=short: -h**

Displays syntax and options for the command. Do not specify other options when you specify the **help** option.

### *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 <code>osecret</code> to <code>nsecret</code> .	<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 **dsmmonitord** 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\_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:

- Use the **errorlogname** option as a command option when you start **dsmwatchd**.
- Use the **errorlogname** option in the `dsm.sys` file to qualify the path and the file name in which to store information about errors that occur during processing. The value of this option overrides the **DSM\_LOG** environment variable.
- Set the environment variable, **DSM\_LOG**. For example:  
`DSM_LOG=/usr/tivoli.tsm/client/hsm/bin/dsmerror.log`
- Restart the **dsmwatchd** daemon. When the **dsmwatchd** daemon starts, it reads the new **errorlogname** option value or the new **DSM\_LOG** environment variable setting.

**On Red Hat Enterprise Linux Version 6 (RHEL6):**

Stop and restart the daemon with the following commands:

```
initctl stop HSM
initctl start HSM
```

**On all other systems besides RHEL6:**

Stop the **dsmwatchd** daemon by using the command `kill -15 dsmwatchd_pid`. The **dsmwatchd** daemon is restarted by the next **init** process.

You can stop the **dsmwatchd** daemon permanently.

**On Red Hat Enterprise Linux Version 6 (RHEL6):**

Stop the daemon with the following command: `initctl stop HSM`.

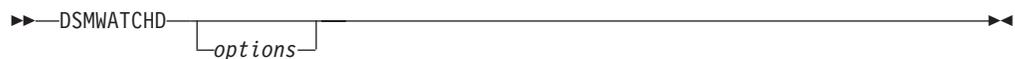
**On all other systems besides RHEL6:**

Remove the **dsmwatchd** entry from the `/etc/inittab` file.

**Note:**

1. For AIX GPFS, set the **DSM\_LOG** environment variable in the `/etc/environment` file accordingly.
2. For Linux x86\_64 GPFS, the `/etc/environment` file does not exist. Use the **errorlogname** option.

## Syntax



## Parameters

*options*

**-ERRORLOGName=***file\_path*

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file `dsm.sys` or `dsm.opt`, or as specified by the **DSM\_LOG** environment variable. If any part of the path you specify does not exist, the HSM client creates it.

**-Help or when optionformat=short: -h**

Displays syntax and options for the command. Do not specify other options when you specify the **help** option.



---

## Chapter 13. HSM GPFS clients Perl scripts

Perl scripts leverage GPFS policy and invoke HSM client commands. You can modify the scripts to suit your environment.

The scripts are valid only for HSM General Parallel File System (GPFS) clients.

The scripts assume that you configured GPFS for integration with the HSM client.

For information about configuring GPFS integration with the HSM client, see the Tivoli field guide *TSM for Space Management for UNIX-GPFS Integration Part I: Policy-driven Threshold Migration* at <https://www.ibm.com/support/docview.wss?uid=swg27018848>.

---

### **dsmMultiServerUpgrade.pl**

Use the **dsmMultiServerUpgrade.pl** script to modify space management of a file system so that the file system can be managed by multiple Tivoli Storage Manager servers.

The script couples all files on a file system with the Tivoli Storage Manager server that manages the migration copies and backup versions. The script calls the **dsmreconcile** command to run a special reconciliation that couples the files with the server.

The script is valid only for HSM GPFS clients. The script is located in the `/opt/tivoli/tsm/client/hsm/multiserver/samples/` directory.

#### **Syntax**

▶—`dsmMultiServerUpgrade.pl`—`file_system_name`—▶

#### **Parameters**

*file\_system\_name*

Specify the file system. All files that are migrated, premigrated, archived, or backed up are coupled with the Tivoli Storage Manager server.

#### **Related tasks**

“Enabling a file system to be managed by multiple Tivoli Storage Manager servers” on page 54

---

## dsmNextServername.pl

Use the `dsmNextServername.pl` script to choose a Tivoli Storage Manager server to manage a file. This script is used for automatic migrations.

If a file is selected for migration and is not coupled with a Tivoli Storage Manager server, this script provides a server name. When the file is migrated, it becomes coupled with the Tivoli Storage Manager server.

The script contains settings that are used to choose a Tivoli Storage Manager server from the list of eligible servers. The script is included with the HSM client and by default uses a round-robin method for choosing servers. You can modify the script to suit your environment. For example, assume that the default server is coupled with many files. When you add a new server, you can exclude the default server until the file management becomes balanced between the new server and the default server. The script is invoked by the GPFS policy engine.

The script is valid only for HSM GPFS clients. The script is located in the `/opt/tivoli/tsm/client/hsm/multiserver/samples/` folder.

### Syntax

```
▶—dsmNextServername.pl—file_system_name—◀
```

### Parameters

*file\_system\_name*

Specify the file system. Typically, the GPFS policy engine provides the file system name.

---

## dsmreconcileGPFS.pl

Use the **dsmreconcileGPFS.pl** script to perform reconciliation in a GPFS environment.

The script uses the GPFS policy engine to determine which files are migrated from the specified file system to a Tivoli Storage Manager server. The script starts **dsmreconcile** with the **orphancheck** and **fileinfo** options.

The two-way orphan check processing identifies orphans on the file system and on the Tivoli Storage Manager server. The two-way check is done in a single pass and the orphan identification process uses parallel processing.

If a file system is managed by multiple Tivoli Storage Manager servers, the script determines which files are migrated to each Tivoli Storage Manager server.

The script creates a list of files that are coupled with a Tivoli Storage Manager server. The script then starts the **dsmreconcile** command:

```
dsmreconcile file_system_name -servername=server_name -orphancheck -fileinfo
-filelist=list_file
```

where *listfile* contains a list of all files on the specified file system *file\_system\_name* that are migrated to the Tivoli Storage Manager server *server\_name*.

The script is valid only for HSM GPFS clients. The script is located in the `/opt/tivoli/tsm/client/hsm/samples/` directory.

## Syntax

```
▶▶ dsmreconcileGPFS servername=server_name file_system_name ▶▶
```

## Parameters

### **servername=server\_name**

Specify the Tivoli Storage Manager server. This option is valid only if the file system is managed by multiple Tivoli Storage Manager servers. All files that are migrated to the specified server from the specified file system are reconciled. Obsolete copies on the server are expired and deleted. If the file system is managed by multiple Tivoli Storage Manager servers, you must specify the **servername** option.

### *file\_system\_name*

Specify the file system to reconcile. The following tasks are performed for the file system:

- The HSM client verifies that valid objects for each local stub file exist in the space management pool of the specified Tivoli Storage Manager server.
- When orphans are identified, their names are recorded in the `.SpaceMan/orphan.stubs` file.

### **Related reference**

“`dsmreconcile`” on page 158

---

## dsmRemoveServer.pl

Use the **dsmRemoveServer.pl** script to reconcile and recall files that are managed by the specified Tivoli Storage Manager server. The recall and reconciliation processes run for the specified file system.

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

1. Recalls all migrated files that are coupled with the Tivoli Storage Manager server that was removed to the specified file system.

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

2. Reconciles the specified file system with the Tivoli Storage Manager server that was removed. The following reconciliation tasks are performed:
  - The HSM client verifies that valid objects for each local stub file exist in the space management pool of the specified Tivoli Storage Manager server.
  - When orphans are identified, their names are recorded in the `.SpaceMan/orphan.stubs` file.
  - The HSM client expires all files in the space management storage pool on the Tivoli Storage Manager server.

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

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

3. Uncouples all files that are coupled with the Tivoli Storage Manager server that was removed.

More than one recall and reconciliation process can run simultaneously.

The script is valid only for HSM GPFS clients. The script is located in the `/opt/tivoli/tsm/client/hsm/multiserver/samples/` directory.

## Syntax

```
▶▶—dsmRemoveServer.pl—server=server_name—file_system_name————▶▶
```

## Parameters

**server=server\_name**

Specify the Tivoli Storage Manager server that will no longer manage the file system.

**file\_system\_name**

Specify the file system that the Tivoli Storage Manager server will no longer manage.

## Related tasks

“Removing a Tivoli Storage Manager server from a multiple-server environment” on page 56

## Related reference

“`dsmreconcileGPFS.pl`” on page 168

---

## Appendix A. Control files in the .SpaceMan directory

When you add space management to your file systems, the HSM client creates control files in a hidden directory named `.SpaceMan`.

Table 29 provides a brief description of the control files and directories in the `.SpaceMan` directory in each space-managed file system. These files are also in the `/etc/adsm/SpaceMan` directory that is created when you install the HSM client.

The HSM client automatically excludes these files from space management. No action on your part is required to ensure that the files remain on the local file systems. The HSM client requires these objects for processing. Do not delete or modify the `.SpaceMan` directory, its contents, or any of the file ownerships and permissions.

*Table 29. Control files stored in the .SpaceMan directory*

Files	Description
<code>ActiveRecallTab</code>	This file contains the active recall table.
<code>config/dmiFSGlobalState</code>	This file contains global and file system information about the Data Management Application Programming Interface (DMAPI) interface.
<code>config/dmiFSGlobalState.pid</code>	This file is a lock file for the <code>dmiFSGlobalState</code> file.
<code>config/DSM.pid</code>	This file is a lock file for the <code>DSMNodeSet</code> and <code>DSMSDRVersion</code> files.
<code>config/DSMNodeSet</code>	This file stores the HSM configuration.
<code>config/DSMSDRVersion</code>	This file stores the number of changes in the HSM configuration.
<code>dmiFSState</code>	This file stores information about the file system.
<code>dsmmigfstab</code>	If this file exists, it is from an older version of the HSM client. It is needed for transition to current version.
<code>hsmfsconfig.pid</code>	This file is a lock file for the local <code>hsmfsconfig.xml</code> file.
<code>hsmfsconfig.xml</code>	This file contains the space management settings for the file system.
<code>logdir/</code>	This directory records information during file migration or recall processing. The information is used to complete any interrupted transactions, such as during a system failure.
<code>metadata/</code>	This directory contains the complete file index (CFI) that was created by the scout daemon. The CFI contains file system information.
<code>orphan.stubs</code>	This file records the orphan files that were identified during reconciliation.
<code>reserved/</code>	This directory contains reserved files for demand migration.
<code>multiserver/BasicRuleSet</code>	This file is a rule set for a file system that is managed by multiple Tivoli Storage Manager servers.
<code>multiserver/serverlist</code>	This file contains the list of Tivoli Storage Manager servers that manage the file system.

Table 29. Control files stored in the `.SpaceMan` directory (continued)

Files	Description
<code>ruleset</code>	This file contains the automatically generated rule set for GPFS.
<code>SDR/</code>	If this directory exists, it is from an older version of the HSM client.
<code>status</code>	This file records space management statistics for the file system.
<code>tapeOptimizedRecall/</code>	This directory contains the recall list files that optimize tape access.

For HSM storage pool support, all files located in the `.SpaceMan` directory are placed in the same storage pool. EXCLUDE rules prevent other migration rules from moving those files to a different pool.

To ensure that other migration rules do not move files into the `.SpaceMan` directory, you can specify the following rules:

```
RULE 'TSM_EXCL_DOTSPACEMAN' EXCLUDE WHERE PATH_NAME LIKE '%/.SpaceMan/%'
```

If files were moved to an unwanted location, the following rule moves all files into one storage pool:

```
RULE 'TSM_MIGR_DOTSPACEMAN' MIGRATE TO POOL 'target_pool'
WHERE PATH_NAME LIKE '%/.SpaceMan/%'
```

where `target_pool` might be "system" or the custom default pool if the system pool is dedicated only to metadata.

#### Related concepts

Chapter 8, "Reconciling file systems," on page 85

"Managing a file system with multiple Tivoli Storage Manager servers" on page 53

"List files for optimized tape recalls" on page 81

## Space requirement for HSM control files

The Space Management client creates control files which require free space on the system. Without sufficient free space, HSM cannot run properly.

Most files in the `.SpaceMan` directory occupy approximately 1 GB of space. In addition to this basic requirement, there are some files that can occupy significantly more space.

#### `.SpaceMan/metadata`

If the scout daemon is running, the daemon creates the complete file index (CFI). The scout daemon runs when you set `hsmdisableautomigdaemons=no` (the default value). The CFI requires approximately 1 KB per file system block or file. For example, if the file system size is 1 TB, and the block size is 1 MB, there are 1,000,000 blocks on the file system. Multiply 1,000,000 blocks by 1 KB per block to yield a CFI size of 1,000,000,000 (1 GB).

You can set the size of the CFI by specifying a maximum number of files to be space managed. Use the `maxfiles` option with the `dsmmigfs update` command. For example, if the file system block size is 1 KB and you specify `maxfiles=1000`, the CFI size is  $1000 \times 1 \text{ KB} = 1 \text{ MB}$ . In this example

| the CFI is only 1 MB, regardless of the total size of the file system. The  
| value of the **maxfiles** option must be greater than or equal to the number  
| of files in the file system.

| **.SpaceMan/reserved**

| When you add space management to a file system, the HSM client creates  
| reserved files. The reserved files require 1 GB of space for each file system  
| that is space managed.

| The total space required can be estimated as 1 MB for base information + 1 GB for  
| reserved files for each space-managed file system. If **hsmdisableautomigdaemons=no**  
| (the default value), you must also add space for the CFI.

| **Related reference**

| “**dsmmigfs add** and **update**” on page 132



---

## Appendix B. Accessibility features for the Tivoli Storage Manager product family

Accessibility features help users who have a disability, such as restricted mobility or limited vision, to use information technology products successfully.

### Accessibility features

The following list includes the major accessibility features in the Tivoli Storage Manager family of products:

- 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

The Tivoli Storage Manager Information Center, and its related publications, are accessibility-enabled. The accessibility features of the information center are described at [http://publib.boulder.ibm.com/infocenter/tsminfo/v6r3/topic/com.ibm.help.ic.doc/iehs36\\_accessibility.html](http://publib.boulder.ibm.com/infocenter/tsminfo/v6r3/topic/com.ibm.help.ic.doc/iehs36_accessibility.html).

### Keyboard navigation

On Windows, the Tivoli Storage Manager product family 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).

On other operating systems, these products follow the operating-system conventions for keyboard navigation and access.

### Vendor software

The Tivoli Storage Manager product family includes certain vendor software that is not covered under the IBM license agreement. IBM makes no representation about the accessibility features of these products. Contact the vendor for the accessibility information about its products.

### IBM and accessibility

See the IBM Human Ability and Accessibility Center for more information about the commitment that IBM has to accessibility.



---

## 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:

*Intellectual Property Licensing  
Legal and Intellectual Property Law  
IBM Japan Ltd  
1623-14, Shimotsuruma, Yamato-shi  
Kanagawa 242-8502 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 want 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 is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

#### COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample

programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows: © (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. \_enter the year or years\_.

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)<sup>®</sup> are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at <http://www.ibm.com/legal/copytrade.shtml>.

Adobe and the Adobe logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Other product and service names might be trademarks of IBM or other companies.



---

## Glossary

This glossary includes terms and definitions for IBM Tivoli Storage Manager and IBM Tivoli Storage FlashCopy Manager products.

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 other 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*.

**bind** To associate all versions of a file with a management class name. See *rebind*.

**bindery**

A database that consists of three system files for a NetWare server. The files contain user IDs and user restrictions.

**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**

An HTTP service that serves the 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 acceptor daemon (CAD)**

See *client acceptor*.

**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.

**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 method**

The method by which a client and server exchange information. See also *Transmission Control Protocol/Internet Protocol*.

**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. 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-programming 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-programming 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.

**extent** The part of a file that is created during the data-deduplication process. Extents are compared with other file extents to identify duplicates.

**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 (ACSL). 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.

**inode** The internal structure that describes the individual files on AIX, UNIX, or Linux systems. An inode contains the node, type, owner, and location of a file.

**inode number**  
A number specifying a particular inode 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** (1) Pertaining to a device, file, or system that is accessed directly from a user system, without the use of a communication line.  
(2) For HSM products, pertaining to the destination of migrated files that are being moved.

**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 backup 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**

**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 Microsoft 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, server definitions, 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.

**MB** See *megabyte*.

**media server**

In a z/OS® environment, a program that provides access to z/OS disk and tape storage for Tivoli Storage Manager servers that run on operating systems other than z/OS.

**megabyte (MB)**

(1) 1 048 576 bytes (2 to the 20th 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 all backed-up versions of a file with a new management class name. For example, a file that has an active backup version is rebound when a later version of the file is backed up with a different management class association. 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 `dsmatrr` 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.

**remote**

(1) Pertaining to a system, program, or device that is accessed through a communication line.

(2) For HSM products, pertaining to the origin of migrated files that are being moved.

**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*.

**system privilege class**

A privilege class that gives an administrator the authority to issue all server commands. See also *privilege class*.

**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.

**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 numeric 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 Microsoft 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 Microsoft 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 Microsoft 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 Microsoft 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 terminal or personal computer 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 48, 171
- .SpaceMan/orphan.stubs 96
- add** 132
- dsmrecall**
  - enhancements for V6.3 xix
- dsmrecall** command 155
- errorlogname**
  - enhancements for V6.3 xix
- filelist** parameter 155
- hsmdisableautomigdaemons**
  - enhancements for V6.3 xix
- hthreshold** 132
- lthreshold** 132
- maxcandidates** 132
- maxfiles** 132
- maxreconcileproc**
  - eliminated in V6.3 xix
- minpartialrecallsize**
  - setting 132
- minstreamfilesize** 132
- pmpercentage** 132
- quota** 132
- stubsizes** 132
- update** 132

## A

- accessibility features 175
- activating space management 137
- add from HSM GUI 44
- add server 136
- adding space management to 45
- adding space management to a file system 132
  - multiple servers 55
- adding to nested file systems 45
- AIX
  - installing HSM 9, 12
- AIX GPFS
  - installing 7
  - requirements 9
  - upgrading HSM client 10
- AIX JFS2
  - hardware requirements 12
  - installing 11
  - requirements 12
  - software requirements 12
- archiving files 67
- archiving migrated files 68
- automatic migration 65
  - by GPFS policy 2
  - dsmNextServername.pl 168
  - file prioritization 62
  - threshold migration 124
- automatic reconciliation 85
- automatically 65
- automignonuse 29

## B

- backing up 67
  - premigrated files 69
- backing up files
  - after migrating 69
  - before migrating 69
- backup and migration 72, 73
- backup-archive client 1
- Bourne and Korn shell 33
- Bourne shell variables 33

## C

- C shell 34
- C shell variables 34
- candidate selection
  - migration 63
- candidates file 171
- candidatesinterval option 25, 95
- checkfororphans option 86, 96
- checkthresholds option 25, 96
- child recall daemon 90
- CIFS 45
- client system options
  - maxrecalldaemons 109
- closed registration
  - HSM UNIX client 19
- clusters on GPFS
  - adding space management 38
- command
  - dsmautomig 124
  - dsmmighelp 118
  - dsmmigquery -options 24
  - dsmmigundelete 70
- command line help
  - displaying 118
- commands 70
  - dsmrecall** 155
  - descriptions 120
  - dmkild 121
  - dsmattr 78, 80, 122
  - dsmautomig 86, 108, 124
  - dsmc 94
  - dsmdf 119, 126
    - incorrect for very young files 61
  - dsmdu 119, 126, 127
    - incorrect for very young files 61
  - dsmls 119, 128, 129, 155
  - dsmmifgs 60
  - dsmmigfs 47, 59, 78, 111, 119, 140, 144
    - adding or updating 132
    - adding space management 44
    - deactivating space management 137
    - globally deactivating space management 138
    - help 139
    - reactivating space management 137

commands (*continued*)

- dsmmigfs (*continued*)
  - removing space management 137
  - takeover 145
- dsmmigfs ADDMultiserver 136
- dsmmigfs command
  - globally deactivating space management 138
  - globally reactivating space management 138
- dsmmigfs rollback 142
- dsmmigfs SDRReset 143
- dsmmighelp 146
- dsmmigquery 30, 146
- dsmmigrate 47, 65, 149
- dsmmigundelete 119, 151
- dsmmonitor 89, 153
- dsmq 92, 154, 161
- dsmrecall 90, 158, 164
- dsmreconcile 86, 87, 96, 158
- dsmrm 92, 161
- dsmrootd 91, 161
- dsmscoutd 162
- dsmscoutd command 164
- dsmsetpw 163
- dsmwatchd 164
  - using 117
- compressalways option 97
- compression 18
- compression option 25, 97
- configuring HSM client 23
- control files
  - HSM 171
- creating stub files for premigrated files 151
- customer support
  - contact xvii

## D

- daemon
  - dsmmonitor 124
  - dsmscout 50
  - dsmscoutd 63
- daemons
  - dsmmigrate 144
  - dsmmonitor 153, 158
  - dsmrecall 144
  - dsmrecall 158
  - dsmrootd 161
  - dsmscoutd 162
  - dsmwatchd 144, 164
  - dsmwatchd daemon 43
  - overview 89
  - scout 90
  - stopping 92
- deactivating
  - space management 58
- deactivating space management 137
  - globally 138
- default size 52

- defaultserver option 25, 28, 98
- defined 3
- definitions 181
- deleted stub files 70
- deleted stub files, restoring 70, 151
- demand migration 2, 62
- disability 175
- disable failover 143
- disaster recovery 75
- disk, restoring 74
- displaying 24
  - space management information 127
  - space management settings 140
  - space usage information 127
- displaying online command help 118
- displaying option information 24
- DMAPI 43
- dmkilld command 121
- dms.sys
  - reconcileinterval 89
- DSM\_CONFIG 33
- DSM\_CONFIG environment variable 33
- DSM\_DIR environment variable 33
- DSM\_LOG environment variable 33
- dsm.opt 95
  - editing 24, 28
- dsm.sys 63, 95
  - editing 24
  - maxcandprocs 63
  - maxrecalldaemons 90
  - minrecalldaemons 90
  - options 25
- dsmattr command 78, 80, 122
- dsmautomig command 65, 108, 124
- dsmc command 94
- dsmdf command 119, 126
- dsmdu command 119, 126, 127
- dsmerror.log 33
- dsmpls command 119, 128, 129, 155
- dsmmifgs
  - removing space management 60
- dsmmigfs
  - command 145
  - reactivate space management 59
  - deactivating space management 137
  - GPFS 136
  - GPFS failover 143
  - help 139
  - query 140
  - reactivating space management 137
- dsmmigfs command 44, 47, 78, 111, 119, 142, 144
  - adding space management 132
  - updating space management settings 132
- dsmmigfs DISABLEFailover 143
- dsmmigfs ENABLEFailover 143
- dsmmigfs QUERYMultiserver 136
- dsmmigfs REMOVEMultiserver 136
- dsmmigfs SDRReset 143
- dsmmighelp command 118, 146
- dsmmigquery
  - shared options 148
- dsmmigquery -options 24
- dsmmigquery command 30, 146
- dsmmigrate command 47, 65, 149
- dsmmigundelete 70

- dsmmigundelete command 70, 119, 151
- dsmmonitor command 90, 153
  - space monitor daemon 89
- dsmmonitor daemon 124, 158
- dsmMultiServerUpgrade.pl 167
  - enabling multiple-server support 54
- dsmNextServername.pl 168
- DSMNodeset file 13
- dsmq command 92, 154, 161
- dsmrecauld command 90, 158, 164
- dsmreconcile
  - command 87
  - reconcileinterval 4
- dsmreconcile command 86, 96, 158
- dsmreconcileGPFS.pl 168
- dsmRemoveServer.pl 169
- dsmrm command 92, 161
- dsmrootd command 91, 161
- dsmssched.log 94
- dsmsscouthd
  - daemon 63
- dsmsscouthd command 162, 164
- dsmsscouthd daemon 26, 95
- DSMSDRVersion file 13
- dsmsetpw command 163
- dsmwatchd command 164

## E

- editing
  - dsm.opt 24
  - dsm.sys 24
- education
  - see Tivoli technical training xiv
- enable failover 143
- enablelanfree 34
- encrypted file system 132
- enhancements for V6.3 xix
- entire system 75
- environment 39
- environment variables
  - DSM\_CONFIG 33
  - DSM\_DIR 33
  - DSM\_LOG 33
- errorlogname option 98
- errorprog option 25, 99
- exclude
  - files (LAN-free) 35
- exclude list
  - edit 32
- exclude options 31
- exclude.compression 97
- excluding files
  - from backup 31
  - from migration 31
  - from space management 31
- expiration 4
  - file 68
- exporting file systems 45
- exporting files using
  - NFS 45

## F

- failover
  - watch daemon 91

- FastConnect 45
- file
  - hsmfsconfig.xml 47
- file list 129
- file size
  - minimum partial 53
- file systems 43, 67, 72
  - adding space management multiple servers 55
  - backing up 67
  - enabling multiple-server support 54
  - from multiple servers 74
  - reconcile 4
  - reconciling 85
  - restoring 67
    - from multiple servers 74
- files 67
  - archiving 67
  - backing up 67
  - migrating 67
  - restoring 67
  - retrieving 67
- fixes, obtaining xvi
- for Linux x86\_64 GPFS 15
- free space requirement
  - HSM control files 172

## G

- globally deactivating space management 138
- globally reactivating space management 138
- glossary 181
- GPFS integration
  - enhancements for V6.3 xix
- GPFS node
  - re-configuring hostname or GPFS node number 46
- GPFS policy-driven automatic migration 2

## H

- HACMP
  - limitations 36
- HACMP for AIX JFS2
  - automatic migration 35
  - enable failover 35
  - to source nodes in cluster 35
- hardware requirements
  - for AIX GPFS 9
  - for AIX JFS2 12
- help 118
- high threshold 49
  - setting 132
- HSM 171
- hsm agent
  - verify running 20
- HSM client
  - installing 5
  - open registration on TSM 18
- HSM client overview 1
- HSM control files 171
  - space requirement 172

- HSM GPFS clients
  - scripts 167
- HSM GUI 20
- HSM UNIX client
  - closed registration on TSM 19
- hsmagent.opt 40
- hsmdisableautomigdaemons
  - enabling multiple-server support 54
- hsmdisableautomigdaemons option 99
- hsmdistributedrecall option 100
- hsmenableimmediatemigrate option 101
- hsmextobjidattr option 25, 101
- hsmfsconfig.xml 43
- hsmfsconfig.xml file 48
- hsmgroupedmigrate option 28, 102
- hsmlogeventflags option 25, 103
- hsmlogmax option 25, 103
- hsmlogname option 25, 104
- hsmlogretention option 25, 105
- hsmlogsampleinterval option 25, 106
- hsmmaxrecalltapedrives option 106
- hsmmultiserver
  - enabling multiple-server support 54
- hsmmultiserver option 107

## I

- IBM Publications Center xi
- IBM Support Assistant xvi
- inexcl option 25, 108
  - dsm.sys 31
- include
  - files (LAN-free) 35
- include list
  - edit 32
- include options 31
- include-exclude
  - file 23
- include.compression 97
- including files
  - for backup 31
  - for migration 31
- inodes 126
- installation directory 5
- installing 13
  - Linux x86\_64 GPFS 15
    - prerequisites 5
    - requirements 5
- installing HSM 5, 13
  - AIX 9, 12
    - AIX JFS2 11
      - on AIX GPFS 7
- Internet, searching for problem
  - resolution xv, xvi

## J

- JFS2 EFS 132
- journal-based backup
  - limitation with GPFS 5

## K

- kernelmessages option 25
- keyboard 175
- knowledge bases, searching xv

- Korn shell variables 33

## L

- LAN-free 32, 34
  - prerequisites 34
- LAN-free configuration
  - limitation
    - multiple servers 57
- LAN-free data transfer 34
- lanfreecommmethod 34
- lanfreetcpport 34
- limitations 5
  - multiple servers 57
- Linux GPFS
  - requirements 15
  - upgrading HSM client 16
- Linux x86\_64 GPFS 13, 15
- list file
  - format 155
- list of files
  - displaying 146
- log files
  - dsmerror.log 33
- logdir directory 171
- logs
  - dsm sched 94
- low threshold 49
  - setting 132

## M

- manage recovery 143
- management 43
- management class
  - assign to files 29
  - default 29
  - displaying 30
- manually 65
- maxcandprocs option 25, 108
- maxmigrators option 25, 108
- maxrecalldaemons
  - option 90
- maxrecalldaemons option 25, 109
- maxthresholdproc option 25, 110
- migdestination 29
- migfileexpiration
  - option 85
- migfileexpiration option 25, 68, 86, 110, 113
- migrated
  - orphans 85
- migrated file
  - setting a recall mode 122
- migrated files
  - recalling 77
- migrateserver option 25, 28, 98, 111
- migrating
  - archived files 68
    - demand 2
    - overview 2
    - recalling files 3
    - selective 2
    - threshold 2
- migrating files
  - archiving 68

- migrating files (*continued*)

- automatic 62
- automatically 124
- demand 62
- selectively 149
- threshold 62
- threshold migration 65
- migration 49
  - automatic 61
    - dsmNextServername.pl 168
  - candidate selection 63
  - eligibility criteria 61
  - overview 61
  - premigration 64
  - selective 61
  - selective migration 149
  - threshold migration 124
- migration candidates list 86, 158
  - displaying 146
- migrequiresbkup 29
- minimum file size 49
- minimum file size for migration 49
- minimum size 52
- minmigfilesize option 25, 111
- minpartialrecallsizes option 78
- minrecalldaemons
  - option 90
- minrecalldaemons option 25, 112
- multiple server
  - dsmreconcileGPFS.pl 168
- multiple servers
  - adding space management 55
  - dsmMultiServerUpgrade.pl 167
  - dsmNextServername.pl 168
  - dsmRemoveServer.pl 169
  - enabling multiple-server support 54
  - limitations 57
  - node replication 57
  - reconciliation 87
  - removing a server 56
  - restoring file systems 74
  - scripts 167

## N

- nested file system
  - mounting parent first 45
- NFS servers 45
- no query restore
  - restoremigstate** option 113
- restoring files 70
- node replication
  - multiple-server environment 57
- normal recall
  - mode 78
  - setting for a migrated file 122

## O

- obsolete files
  - reconciliation 4
- online help
  - displaying for commands 118, 146
- open registration
  - HSM client 18
- operating system 75

- optimized tape recall 3, 80
- optional setup 29
- optionformat option 25, 28, 112
- options 24, 34, 95
  - candidatesinterval 25, 90, 95
  - checkfororphans 86, 87, 96
  - checkthresholds 25, 96
  - compressalways 97
  - compression 25, 97
  - defaultserver 25, 28, 98
  - enablelanfree 34
  - errorlogname 98
  - errorprog 25, 87, 99
  - exclude 31
  - for scheduling services 93
  - hsmdisableautomigdaemons 99
  - hsmdistributedrecall 100
  - hsmenableimmediatemigrate 101
  - hsmextobjidattr 25, 101
  - hsmgroupedmigrate 28, 102
  - hsmlogeventflags 25, 103
  - hsmlogmax 25, 103
  - hsmlogname 25, 104
  - hsmlogretention 25, 105
  - hsmlogsampleinterval 25, 106
  - hsmmaxrecalltapedrives 106
  - hsmmultiserver 107
  - inclxcl 25, 108
  - include 31
  - kernelmessages 25
  - lanfreecommmethod 34
  - lanfreetcport 34
  - maxcandprocs 25, 108
  - maxmigrators 25, 108
  - maxrecalldaemons 25, 90, 109
  - maxthresholdproc 25, 110
  - migfileexpiration 25, 68, 85, 86, 110, 113
  - migrateserver 25, 28, 98, 111
  - minmigfilesize 25, 111
  - minrecalldaemons 25, 90, 112
  - optionformat 25, 28, 112
  - passwordaccess 19
  - reconcileinterval 25, 85, 93, 113
  - reconciliation
    - tasks 85
  - restoremigrate 68
  - restoremigstate 25, 28, 113
  - restoreMigState 86
  - rollback 142
  - schedlogname 94
  - servelastaccessdate 61
  - setting space management 23
  - standard format 117
  - streamseq 115
- ordered recall list
  - displaying 146
- orphan stub files 87
- orphan stubs 87
- orphan.stubs file 85, 171
- orphans 85
- overview 5

## P

- packages
  - AIX JFS2 12

- parent
  - mounting before nested file systems 45
- partial file
  - minimum size 53
- partial file recall 3, 78
- partial recall mode
  - setting for a migrated file 122
- Passport Advantage xvii
- password 163
  - encryption 19
  - generate 19
  - passwordaccess 19
  - TSM 19
- perl scripts 167
  - dsmMultiServerUpgrade.pl 167
  - dsmNextServername.pl 168
  - dsmreconcileGPFS.pl 168
  - dsmRemoveServer.pl 169
- policies
  - displaying 29, 30
- policy reconciliation
  - GPFS policy 87
- port mapper service 91
- portmap service 91
- premigrated files
  - orphaned 85
- premigrating
  - files 69
- premigrating files
  - archiving 68
- premigration 64, 65
- premigration percentage 50, 64
- premigration setting 132
- premigrdm.dir file 171
- premigrdm.pag file 171
- prerequisites 34
  - installing 5
  - LAN-free 34
- problem determination
  - describing problem for IBM Software Support xviii
  - determining business impact for IBM Software Support xvii
  - submitting a problem to IBM Software xviii
- publications
  - download xi
  - order xi
  - search xi
  - Tivoli Storage Manager xii

## Q

- query server 136
- quota 51
  - setting 132

## R

- reactivating space management 59, 137
  - globally 138
- read-only recall 78, 79
- recall
  - daemon 90
  - HACMP 35

- recall (*continued*)
  - tape-optimized recall 80
- recall daemon 27
  - dmkill command 121
  - starting 158
- recall ID 154
- recall mode
  - dsmattr command 80
  - HSM determines 80
  - normal 122
  - setting for a migrated file 122
- recall process
  - displaying status of 154
- recalling files 78, 80
  - canceling 161
  - migrated files 3
  - normal 78
  - removing from queue 161
  - selective recall 77
  - selectively 77, 155
  - streaming 79
  - transparent recall 77
- recallmode
  - option 78
- recommendations
  - not adding space management 44
- reconcile file systems 4
- reconcileinterval
  - dms.sys 89
  - option 85
- reconcileinterval option 25, 93, 113
- reconciliation
  - automatic 85
  - dsmreconcileGPFS.pl 168
  - GPFS policy 87
  - multiple servers 87
  - policy-based 158
- reconciling
  - manual 85
  - manually 86, 158
- reconciling file systems 158
- recreating deleted stub files 70, 151
- remotely mounted file systems
  - no HSM support 5
- remove server 136
- removing
  - space management 137
- removing a server
  - multiple-server environment 56
- requirements
  - AIX GPFS 9
  - AIX JFS2 12
  - installing 5
  - Linux GPFS 15
- reset in the SDR 143
- resolving 87
- Responsiveness Service 91
- restarting space management
  - daemons 144
- restore migrated files 70
- restoremigrate 68
- restoremigstate
  - command 87
- restoremigstate option 25, 28, 113
- restoreMigState option 86
- restoring 67, 70, 72, 74, 75
  - a disk 74

- restoring *(continued)*
  - deleted stub files 151
  - premigrated files 69
- restrictions 5
- retrieving
  - archived files 68
- retrieving files 67
- rollback option 142
- root daemon 91
  - starting 161
- root user tasks 153
  - creating stub files for premigrated files 151
  - recreating deleted stub files 151
  - starting a root daemon 161
  - starting a scout daemon 162

## S

- SAMBA 45
- sample script 36
- SAN
  - LAN-free 34
- schedlogname option 94
- scheduled services
  - displaying 94
  - displaying completed services 94
- scheduler
  - starting 93
- scheduling
  - backup-archive services 93
  - space management services 93
- scheduling services
  - options 93
- scout daemon 90
  - starting 162
- scripts
  - for HSM GPFS clients 167
  - perl 167
- selected files
  - migrating 149
  - recalling 155
- selecting HSM options 23
- selective migration 2, 65
  - performing 149
- selective recall 3, 77, 155
- servelastaccessdate option 61
- server is different 73
- server is the same 72
- server name encryption limitation
  - multiple servers 57
- server node replication limitation
  - multiple servers 57
- services
  - scheduling 93
- setting
  - quota 51
- setting environment variables 33
- setting space management options 23
- setting thresholds for 49
- setting up 23
- settings
  - maximum file size 53
  - minimum partial file size 53
  - minimum streaming file size 52

- settings *(continued)*
  - premigration percentage 50
  - updating from HSM GUI 47
- setup tasks 29
- shell scripts 1
- software requirements 15
  - for AIX GPFS 9
  - for AIX JFS2 12
- software support
  - describing problem for IBM Software Support xviii
  - determining business impact for IBM Software Support xvii
  - submitting a problem xviii
- Software Support
  - contact xvii
- space management 23, 44, 45
  - adding to a file system 132
  - agent 40
  - clusters on GPFS 38
  - files
    - candidates file 171
    - logdir directory 171
    - orphan.stubs file 171
    - premgrdb.dir file 171
    - premgrdb.pag file 171
    - status file 171
  - installing 5
  - maintaining 132
  - removing 60, 137
  - settings 47
  - updating a file system 132
- space management settings 132
  - maxcandidates** 132
  - displaying 140
  - high threshold 132
  - premigration setting 132
  - quota 132
  - stub file size 132
- space monitor daemon 26, 89, 96
  - starting 153
- space requirement
  - HSM control files 172
- space usage information
  - displaying 127
- space usage, displaying 99, 100, 102, 103, 105, 106, 107, 119, 126
- SpaceMan directory
  - setting low threshold 49
- spacemgtechnique 29
- standard option format 117
- start HACMP 36
- starting
  - client scheduler 93
  - space management daemons 144
- status file 171
- stop HACMP 36
- stopping space management daemons 92, 144
- storage pool 1, 90
  - dsmautomig command 124
- storage pools 65
- stream file size 52
- streaming recall 3
- streaming recall mode 79
- streamseq option 115

- stub file size
  - setting 132
- stub files
  - default size by platform 52
  - leader data 52
  - maximum size 52
  - minimum size 52
  - orphan 87
  - recreating 151
  - recreating deleted 119
  - restoring 70
  - setting the size 52
- summary of changes V6.3 xix
- support contract xvii
- support information xiv
- support subscription xvii
- supported file systems xix
- synchronizing client-server files 85
  - manual reconciliation 158

## T

- takeover filespec 145
- tape recall, optimized 80
- tasks 29
- threshold migration 2, 62
  - start manually 65
  - starting manually 124
- thresholds
  - migration 49
- Tivoli Storage Manager password 163
- Tivoli Storage Manager servers, multiple 66
- Tivoli technical training xiv
- training, Tivoli technical xiv
- transfer HSM management of file system 142, 145
- transparent recall 3, 77
- truncate
  - migrated files 3
- two-way orphan check 158

## U

- uninstalling
  - from AIX GPFS 11
  - uninstalling from Linux 17
- update files in SDR 143
- updating space management settings 132
- upgrade
  - multiple servers
    - dsmMultiServerUpgrade.pl 167
- upgrading HSM client
  - AIX GPFS 10
  - Linux GPFS 16
- using commands 117

## V

- variables 33, 34
- variables, environment 33
- verifying
  - hsm agent running 20
- Version 6.3 changes xix

## W

watch daemon 43, 91

workstation

    registering with a server 18

WPAR 46

    use HSM only in global WPAR 5





Product Number: 5608-E12

Printed in USA

SC23-9794-03

