IBM Tivoli Storage Manager for Space Management Version 7.1

User's Guide



IBM Tivoli Storage Manager for Space Management Version 7.1

User's Guide



Note:

Before using this information and the product it supports, read the information in "Notices" on page 195.

First edition (December 2013)

This edition applies to version 7, release 1, modification 0 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.

© Copyright IBM Corporation 1995, 2013.

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

Contents

Figures
Tables
Preface xi Who should read this publication xi Publications xi Tivoli Storage Manager publications xi Conventions used in this publication xiiii
Tivoli Storage Manager for Space Management Version 7.1 updates xv
Chapter 1. Tivoli Storage Manager for Space Management client overview 1 Migrating files overview
Chapter 2. Installing the Tivoli Storage Manager for Space Management client 7
General installation prerequisites and considerations 7 Installation overview of Tivoli Storage Manager for Space Management client for AIX GPFS systems 9 Environment requirements for the Tivoli Storage
Manager for Space Management client for AIX GPFS systems
Management client for AIX systems
Uninstalling the Tivoli Storage Manager for Space Management client for AIX GPFS systems . 14 Installation overview of the Tivoli Storage Manager
for Space Management client for AIX JFS2 systems . 14 Environment requirements for the Tivoli Storage Manager for Space Management client for AIX
JFS2 systems
Installation overview for the Tivoli Storage Manager for Space Management client for Linux x86_64 GPFS
Systems
Installing the Tivoli Storage Manager for Space Management client for Linux x86_64 GPFS
Upgrading the Tivoli Storage Manager for Space Management client for Linux x86_64 GPFS
systems

Uninstalling the Tivoli St	ora	ige	Ma	ana	gei	fo	r	DEC	-	
Space Management clien	t fc	or L	ını	1X)	(86	_64	G	PFS	>	
systems										20
Registering your workstatio	n v	vith	ı a	Tiv	oli	Sto	ora	ge		
Manager server										21
Open registration										21
Closed registration										21
Managing your password										22
The Tivoli Storage Manager	fo	r Sp	bac	e N	lan	age	em	ent		
client GUI										22
Verifying that the HSM a	igei	nt i	s r	unr	ning	z .				23
Using the Tivoli Storage	Мa	inaș	ger	for	r Si	bac	е			
Management client GUI										24
0										

Chapter 3. Configuring the Tivoli Storage Manager for Space

Management client		25
Displaying Tivoli Storage Manager for Space		
Management client options		. 26
Editing the options files by using the		
backup-archive client		. 26
Editing the options file dsm.sys		. 27
Editing the options file dsm.opt		. 28
Configuring the Tivoli Storage Manager for Space	ć	
Management client to connect to a secondary Tive	oli	
Storage Manager server		. 29
Optional setup features		. 30
Assignment of management classes to files .		. 30
Displaying management class information .		. 32
Incude-exclude file options		. 32
Setting environment variables		. 34
LAN-free data transfer for HSM		. 36
HSM for AIX GPFS and Linux x86_64 GPFS		
clusters		. 37
HSM space management agent configuration		. 38
Logs for HSM activity and error messages .		. 39

Chapter 4. Adding and configuring

space management for a file system 41
Adding space management to file systems 42
Adding space management to nested file systems . 43 Mounting the parent file system before nested
file system
Adding space management to an exported file
system
Removing or reconfiguring a space-managing node
in a GPFS cluster
Adding space management to workload partitions
on AIX V6.1 and V7.1 operating systems 44
Updating settings from the command line 45
Updating settings from the Tivoli Storage Manager
for Space Management client GUI
Settings that control the space usage of a file system 46
Minimum migration file size
Migration threshold percentages

Premigration percentage	49
Quotas	50
Stub file size	50
Minimum stream file size.	51
Minimum partial file size.	51
Maximum number of files	51
Managing a file system with multiple Tivoli Storage	
Manager servers.	52
Enabling a file system to be managed by multiple	
Tivoli Storage Manager servers	52
Adding space management to a file system in a	
multiple-server environment	54
Removing a Tivoli Storage Manager server from	01
a multiple-server environment	54
I imitations for multiple-server environments	55
Replicating podes in a multiple-server	00
anvironment	56
Descripting space management	57
Reactivating space management	57
Reactivating space management	57
Removing space management	38
Obserter 5 Minusting files	~1
Chapter 5. Migrating files	61
Migration types	61
File migration eligibility	62
File premigration	63
Automatic file migration	63
Candidate selection for automatic migration	64
Manually starting threshold migration	65
Selective file migration	66
Migrating selectively using the domnigrate	
wingrating selectively using the dominigrate	
command	66
command	66
command . </td <td>66 67</td>	66 67
wind the generative structure is a single file system to two or more Tivoli Storage Manager servers	66 67
Wightamig screenvery using the distinguistic command	66 67
command	66 67 69
wind it is selectively using the usining file command	66 67 69
wind it in a single second y using the distinguistic command	66 67 69 70
Wighting screenvery using the distinguite command	66 67 69 70 70
Wighting screenvery using the usiningrate command	66 67 69 70 70 70 70
Wightang screenvery using the usiningrate command	66 67 69 70 70 70 70 71
Wightaming screenvery using the distinuing factoring command	66 67 69 70 70 70 71 71
Wightaming screenvery using the distinlighted command	66 67 69 70 70 70 70 71 71
Wightaming screenvery using the distinlighted command	66 67 69 70 70 70 71 71 71
Wighting screenvery using the usiningrate command	66 67 69 70 70 70 71 71 71 72 73
Wighting screenvery using the usiningrate command	66 67 69 70 70 70 71 71 71 72 73
Wightling screenvery using the usiningrate command	66 67 69 70 70 70 71 71 71 72 73 74
Wightling screenvery using the usiningrate command	66 67 69 70 70 70 71 71 71 72 73 74
Wightling screenvery using the usining file command	66 67 69 70 70 70 70 71 71 71 72 73 74
Wightling screenvery using the usining the usining the command	 66 67 69 70 70 70 71 71 72 73 74 74
Wightling screenvery using the usining the usining the command	 66 67 69 70 70 70 71 71 72 73 74 74
Wightling screenvery using the usining the usining the command	66 67 69 70 70 71 71 71 72 73 74 74
Wightling screenvery using the usinfigure command	 66 67 69 70 70 70 71 71 72 73 74 74 75 76
Wightling screenvery using the usinfigure command	 66 67 69 70 70 70 71 71 72 73 74 74 75 76 76 76 76 76
Wightling screenvery using the usinfigure command	66 67 69 70 70 70 71 71 72 73 74 74 75 76 76 76
Wightling screenvery using the usinfigure command	66 67 69 70 70 70 71 71 71 72 73 74 74 74 75 76 76 76 77
Wightling selectively using the usinfigure command	66 67 69 70 70 71 71 71 72 73 74 74 74 74 75 76 76 76 77
Wightling screenvery using the usinfigure command	66 67 69 70 70 71 71 71 72 73 74 74 74 74 75 76 76 76 77 77

Streaming recall mode					
Recalling files with optimization for tape 84					
Chapter 8. Reconciling file systems 85					

			. 85
			. 86
			. 87
			. 88
			. 88
	· · · · · ·	· · · ·	· · · · · · · · · ·

Chapter 9. Space management

deemene	00
	. 89
The space monitor daemon	. 89
The recall daemon	. 90
The scout daemon	. 90
The watch daemon	. 91
Stopping the space management daemons	. 91
Chapter 10. Scheduling services	. 93
Scheduling options	. 93
Starting the client scheduler	. 93
Displaying scheduled services information	. 94
1 7 0	
Chapter 11. Options files reference	. 95
candidatesinterval	95
checkfororphans	. 20
checkthresholds	. 96
compression	. 97
defaultserver	. 98
errorlogname	. 98
errorprog	. 90
hsmdisableautomiddaemons	
hsmaistrihutedrecall	100
hsmonahleimmediatemiorate	101
hsmeventdestrov	101
hsmevtohiidattr	102
hsmarounodmiarato	102
hsmloreventflars	102
hsmlogevenerings	104
	105
hsmlogratention	105
hsmlogsamnleinterval	106
hsmmaxrecalltanedrives	107
hsmmulticarver	107
inclexcl	108
maxcandprocs	108
maxmigrators	100
maximigrations	109
maxtecalluaemons	. 107
martineshoruproc	. 110
migrifecapitation	. 110
migrateserver	. 111
minimigrifesize	. 112
reconcileinterval	. 112
reconcile interval	. 113
icoloiciliyolale	. 113

Normal recall mode

Chapter 12. HSM for UNIX and Linux

client cor	nn	nai	nd	re	fe	rer	nce	Э.						•	115
Standard op	otio	n fo	orm	nats	5.										115
Displaying	com	nma	and	l-lir	ne I	help	5								116
Displaying	file	an	d fi	ile s	sys	terr	i in	for	ma	tio	n				116
Client return	n co	ode	s												117
HSM comm	and	d sı	ımı	mai	ry										118
dmkilld .															120
dsmattr .															120
dsmautomig															123
dsmdf															125
dsmdu															126
dsmls															128
dsmmigfs ad	d ar	nd	upo	lat	е										131
dsmmigfs ad	dmu	lti	ise	rve	r, (que	rym	u11	tis	erv	er,	an	d		
removemulti	ser	ve	r												135
dsmmigfs de	act	iva	te,	, re	ac	tiv	ate	, ai	nd	rem	ov	е			137
dsmmigfs gl	oba	1de	ac	tiv	ate	e ar	nd g	g1 o	bal	rea	act	iva	ite		138
dsmmigfs he	1p			•										•	139
dsmmigfs qu	ery			•	•						•			•	139
dsmmigfs ro	11b	ack	ς.	•										•	142
dsmmigfs sd	rre	set	t, e	nab	le	fai	lov	er,	an	d					
disablefail	ove	er	•		•	•				•	•				143
dsmmigfs st	op,	sta	art	, ar	nd	res	tar	٠t			•			•	144
dsmmigfs ta	keo	ver	۰.		•	•				•	•				145
dsmmighelp				•										•	146
dsmmigquery	′ .	•	•	•	•	•	•	•	•		•	•		•	147
dsmmigrate		•	•	•	•	•	•	•	•		•	•		•	149
dsmmigundel	ete	2.	•	•	•	•	•	•	•		•	•		•	152
dsmmonitord	Ι.	•	•	•	•	•	•	•	•	•	•	•	•	•	154
dsmq	•	•	•	•	•	•	•	•	•	•	•	•	•	•	155
dsmrecall		•	•	•	•	•	•	•	•		•	•		•	156
dsmrecalld	•	•	•	•	•	•	•	•	•	•	•	•	•	•	159
dsmreconcil	е	•	•	•	•	•	•	•	•	•	•	•	•	•	160
dsmrm		•	•	•	•	•	•	•	•		•	•		•	162
dsmscoutd		•	•	•	•	•	•	•	•		•	•		•	163
dsmsetpw .		•	•	•	•	•	•	•	•	•	•	•	•	•	164
dsmwatchd				•	•	•	•	·	·	•	•	•	•	•	165
Chapter 1	3.	H	SN	1 G	iΡ	FS	cl	ie	nts	s P	er	I			

. ecrinte

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

Chapter 14. Troubleshooting the Tivoli Storage Manager for Space

Management client		173
Client trace classes for the Tivoli Storage Mana	ager	
for Space Management client		. 173
Enabling a Space Management client trace		. 174
Enabling daemon traces		. 175
Tracing the watch daemon		. 175
Tracing the recall daemon		. 176
Tracing the scout daemon		. 176
Tracing the monitor daemon		. 177
Log files to aid troubleshooting		. 177
Commands for displaying HSM information .		. 177
Commands for displaying GPFS information .		. 179

Commands for displaying operating syste	em			
information			. 12	79
Other tools for collecting information .			. 18	80
HSM common problems and solutions .			. 18	80

Appendix A. Control files in the

.SpaceMan directory	•	•						183
Space requirement for HSM	CO	ntro	ol	files	5			184

Appendix B. Tivoli support

information	187
Communities and other learning resources	. 187
Searching knowledge bases	. 189
Searching the Internet	. 189
Using IBM Support Assistant	. 189
Finding product fixes.	. 190
Receiving notification of product fixes	. 190
Contacting IBM Software Support	. 190
Setting up and managing support contracts .	. 190
Determining the business impact	. 191
Describing the problem and gathering	
background information.	. 191
Submitting the problem to IBM Software	
Support	. 191

Appendix C. Accessibility features for

th fa	e mi	Гiv Iy.	oli	i S	ito	ra	ge	M	an	ag	jer	p	roo	du	ct	 . 193
No	oti	ce	s.													 . 195
Tra	nde	ma	rks													. 197
Pri	va	cy]	pol	icy	co	nsi	der	ati	ons		•			•		. 197
GI	os	sa	ary													 . 199
А																. 199
В																. 201
С																. 202
D																. 203
Е																. 205
F																. 206
G																. 206
Η																. 207
I.																. 208
J.																. 208
K																. 208
L																. 209
Μ																. 210
Ν																. 211
Ο																. 212
Р																. 212
Q																. 213
R																. 214
S																. 215
Т																. 218
U																. 218
V																. 219
W																. 220
In	de	х.														 . 221

Figures

1.	XML configuration file				. 38
2.	Ordered list file generation				. 82

Tables

1.	Server publications
2.	Storage agent publications
3.	Backup-archive client publications xii
4.	Data protection publications
5.	Troubleshooting and performance tuning
	publications
6.	Tivoli Storage Manager for Space Management
	client for AIX GPFS systems installation
	packages
7.	Tivoli Storage Manager for Space Management
	client for AIX JFS2 systems installation
	packages
8.	Linux GPFS available packages
9.	Package names and descriptions
10.	Tivoli Storage Manager for Space Management
	client options in the dsm.opt file
11.	Tivoli Storage Manager for Space Management
	client attributes in a management class 31
12.	Assigning management classes to your files 31
13.	Include and Exclude Statements
14.	Environment Variables
15.	LAN-free data transfer options

16.	HSM agent options
17.	dsmmigundelete command actions without the
	expiring option
18.	dsmmigundelete command actions with the
	expiring option
19.	Tivoli Storage Manager for Space Management
	client determining recall modes
20.	Automatic reconciliation tasks
21.	Option format examples: options without
	values
22.	Option format examples: options with values 116
23.	HSM commands to display file and file
	system information
24.	An explanation of client return codes 117
25.	HSM command summary
26.	Space management for locally and remotely
	mounted file systems
27.	Space management for locally and remotely
	mounted file systems
28.	Trace classes
29.	Common HSM problems and resolutions 180
30.	Control files stored in the .SpaceMan directory 183

Preface

This publication provides the information to install, configure, monitor, and troubleshoot problems with IBM[®] Tivoli[®] Storage Manager for Space Management.

Who should read this publication

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

Publications

Publications for the Tivoli Storage Manager family of products are available online. The 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 across all publications or to download PDF versions of individual publications, go to the Tivoli Storage Manager information center at http://pic.dhe.ibm.com/infocenter/tsminfo/v7r1.

You also can find the Tivoli Storage Manager product family information centers and other information centers that contain official product documentation for current and previous versions of Tivoli products at Tivoli Documentation Central. Tivoli Documentation Central is available at http://www.ibm.com/ developerworks/community/wikis/home/wiki/Tivoli Documentation Central.

Tivoli Storage Manager publications

The following tables list the publications that are included in the Tivoli Storage Manager library.

Publication title	Order number
IBM Tivoli Storage Manager for AIX Installation Guide	GC23-9781
IBM Tivoli Storage Manager for AIX Administrator's Guide	SC23-9769
IBM Tivoli Storage Manager for AIX Administrator's Reference	SC23-9775
IBM Tivoli Storage Manager for HP-UX Installation Guide	GC23-9782
IBM Tivoli Storage Manager for HP-UX Administrator's Guide	SC23-9770
IBM Tivoli Storage Manager for HP-UX Administrator's Reference	SC23-9776
IBM Tivoli Storage Manager for Linux Installation Guide	GC23-9783
IBM Tivoli Storage Manager for Linux Administrator's Guide	SC23-9771
IBM Tivoli Storage Manager for Linux Administrator's Reference	SC23-9777
IBM Tivoli Storage Manager for Oracle Solaris Installation Guide	GC23-9784
IBM Tivoli Storage Manager for Oracle Solaris Administrator's Guide	SC23-9772
IBM Tivoli Storage Manager for Oracle Solaris Administrator's Reference	SC23-9778

Table 1. Server publications

Table 1. Server publications (continued)

Publication title	Order number
IBM Tivoli Storage Manager for Windows Installation Guide	GC23-9785
IBM Tivoli Storage Manager for Windows Administrator's Guide	SC23-9773
IBM Tivoli Storage Manager for Windows Administrator's Reference	SC23-9779
IBM Tivoli Storage Manager for z/OS Media Installation and User's Guide	SC27-4018
IBM Tivoli Storage Manager Upgrade and Migration Guide for V5 Servers	GC27-4017
IBM Tivoli Storage Manager Integration Guide for Tivoli Storage Manager FastBack	SC27-2828

Table 2. Storage agent publications

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

Table 3. Backup-archive client publications

Publication title	Order number
IBM Tivoli Storage Manager for UNIX and Linux: Backup-Archive Clients Installation and User's Guide	SC23-9791
IBM Tivoli Storage Manager for Windows: Backup-Archive Clients Installation and User's Guide	SC23-9792
IBM Tivoli Storage Manager Using the Application Programming Interface	SC23-9793
IBM Tivoli Storage Manager for Space Management for UNIX and Linux: User's Guide	SC23-9794
IBM Tivoli Storage Manager HSM for Windows Administration Guide	SC23-9795

Table 4. Data protection publications

Publication title	Order number
IBM Tivoli Storage Manager for Databases: Data Protection for Microsoft SQL Server Installation and User's Guide	GC27-4010
IBM Tivoli Storage Manager for Databases: Data Protection for Oracle for UNIX and Linux Installation and User's Guide	SC27-4019
IBM Tivoli Storage Manager for Databases: Data Protection for Oracle for Windows Installation and User's Guide	SC27-4020
IBM Tivoli Storage Manager for Mail: Data Protection for Microsoft Exchange Server Installation and User's Guide	GC27-4009

Table 4. Data protection publications (continued)

Publication title	Order number
IBM Tivoli Storage Manager for Mail: Data Protection for Lotus [®] Domino [®] UNIX and Linux Installation and User's Guide	SC27-4021
IBM Tivoli Storage Manager for Mail: Data Protection for Lotus Domino for Windows Installation and User's Guide	SC27-4022
IBM Tivoli Storage Manager for Enterprise Resource Planning: Data Protection for SAP Installation and User's Guide for DB2 [®]	SC33-6341
IBM Tivoli Storage Manager for Enterprise Resource Planning: Data Protection for SAP Installation and User's Guide for Oracle	SC33-6340
IBM Tivoli Storage Manager for Virtual Environments Installation and User's Guide	SC27-2898

Table 5. Troubleshooting	and	performance	tuning	publications
--------------------------	-----	-------------	--------	--------------

Publication title	Order number
IBM Tivoli Storage Manager Problem Determination Guide	GC23-9789
IBM Tivoli Storage Manager Performance Tuning Guide	GC23-9788
IBM Tivoli Storage Manager Client Messages and Application Programming Interface Return Codes	SC27-2878
IBM Tivoli Storage Manager Server Messages and Error Codes	SC27-2877
IBM Tivoli Storage Manager for Mail: Data Protection for Microsoft Exchange Server Messages	GC27-4011
IBM Tivoli Storage Manager for Databases: Data Protection for Microsoft SQL Server Messages	GC27-4012
IBM Tivoli Storage Manager for Databases: Data Protection for Oracle Messages	SC27-4014
IBM Tivoli Storage Manager for Mail: Data Protection for Lotus Domino Messages	SC27-4015
IBM Tivoli Storage Manager for Enterprise Resource Planning: Data Protection for SAP Messages	SC27-4016

Note: IBM System Storage[®] Archive Manager is not releasing a new product. You can find information about IBM System Storage Archive Manager at the Tivoli Storage Manager V6.3 information center (http://pic.dhe.ibm.com/infocenter/tsminfo/v6r3/topic/com.ibm.itsm.nav.doc/c_complydataretention_ovr.html).

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

Example	Description		
dateformat	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 7.1 updates

The following features are new for IBM Tivoli Storage Manager for Space Management Version 7.1. Tivoli Storage Manager for Space Management V7.1 gives you greater control of streaming recall, shows the status of GPFS[™] nodes, and can recall files to resident state. Troubleshooting information is published.

Control of streaming recall is improved

In previous versions, the buffering of streaming recalls was set by using the **streamseq** option in the dsm.sys option file. In version 7.1 and later versions, you set streaming recall attributes by using the **dsmmigfs add** and **dsmmigfs update** commands. The **dsmmigfs add** and **dsmmigfs update** commands have new parameters that set the following attributes:

- The streaming recall timeout value.
- Whether a file is recalled as soon as the stub is read.

For more information about the **dsmmigfs add** and **dsmmigfs update** commands, see "**dsmmigfs add** and **update**" on page 131.

You can query the status of GPFS nodes by using the dsmmigfs query -node command

The new **node** option, which is available with the **dsmmigfs query** command, displays information about GPFS nodes. You can display the output in an easily parsable format by using the **parsable** option.

For more information about the **node** option and the **parsable** option, see "dsmmigfs query" on page 139.

You can specify that recalled files are in resident state

In previous versions, recalling a file did not delete the migrated copy in Tivoli Storage Manager storage. The recalled file was in premigrated state. In version 7.1, the **resident** option specifies that recalled files are in resident state. The migrated copy in storage is expired at the next reconciliation.

For more information about the **resident** option, see "**dsmrecal1**" on page 156.

You can manually configure a connection to a secondary Tivoli Storage Manager

server When the primary Tivoli Storage Manager server for HSM fails, you can connect to a secondary server. The primary Tivoli Storage Manager server for the Tivoli Storage Manager for Space Management client must be configured to replicate client node data. You can only recall files from the secondary server; you cannot migrate files to the secondary server.

For more information about configuring the connection to a secondary Tivoli Storage Manager server, see "Configuring the Tivoli Storage Manager for Space Management client to connect to a secondary Tivoli Storage Manager server" on page 29.

The minimum abbreviation of the **stubsize** option is increased from 2 characters to 3 characters

You use the **stubsize** option to specify the size of stub files that remain on your local file systems when files are migrated to storage. The minimum

abbreviation of the **stubsize** option is increased from **st** to **stu**. The three-character abbreviation distinguishes the **stubsize** option from the **streamseq** option.

For more information about the **stubsize** option, see "**dsmmigfs add** and **update**" on page 131.

Troubleshooting information is added to the Tivoli Storage Manager for Space Management User's Guide

Information is provided that can help you resolve some typical problems. Instructions for gathering troubleshooting information can you help you and IBM representatives resolve problems.

For more information, see Chapter 14, "Troubleshooting the Tivoli Storage Manager for Space Management client," on page 173.

Related concepts:

Chapter 1, "Tivoli Storage Manager for Space Management client overview," on page 1

Chapter 1. Tivoli Storage Manager for Space Management client overview

The IBM Tivoli Storage Manager for Space Management 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.



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

File migration, unlike file backup, does not protect against accidental file deletion, file corruption, or disk failure. Continue to back up your files regardless of whether they 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 or the complete files.

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

The Tivoli Storage Manager for Space Management client functions for threshold migration, demand migration, selective migration, selective and transparent recall includes processing GPFS file systems containing multiple space-managed storage pools.

The Tivoli Storage Manager for Space Management client has both a graphical user interface (the Tivoli Storage Manager for Space Management client 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 Tivoli Storage Manager for Space Management 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.

Related concepts:

Chapter 2, "Installing the Tivoli Storage Manager for Space Management client," on page 7

"File migration eligibility" on page 62

Chapter 6, "Backing up and restoring file systems," on page 69

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 7.1 updates" on page xv

Migrating files overview

The Tivoli Storage Manager for Space Management client provides both automatic and selective migration. Once file migration begins, the Tivoli Storage Manager for Space Management 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 Tivoli Storage Manager for Space Management client provides two types of automatic migration: threshold migration and demand migration.

Threshold migration maintains a specific level of free space on your local file system. When space usage reaches the high threshold that you set for your file system, eligible files are migrated to 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 Tivoli Storage Manager for Space Management 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:

"File migration eligibility" on page 62

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 Tivoli Storage Manager for Space Management client recalls a migrated file.

Normal recall mode recalls a migrated file to its originating file system. The recalled file remains on your local file system. When you close the unmodified file, the copy that currently resides in 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 Tivoli Storage Manager for Space Management client intercepts a read request for a file configured for partial file recall. The Tivoli Storage Manager for Space Management client then calculates which portion of the file to recall based on the offsets contained in the read request. This results in time and disk space savings, since only a portion of the file is recalled.

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

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

Related concepts:

Chapter 7, "Recalling migrated files," on page 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 Tivoli Storage Manager for Space Management client

You install the Tivoli Storage Manager for Space Management 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 Tivoli Storage Manager for Space Management client.

The Tivoli Storage Manager for Space Management 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, "Tivoli Storage Manager for Space Management client overview," on page 1

"Assignment of management classes to files" on page 30

"Registering your workstation with a Tivoli Storage Manager server" on page 21

General installation prerequisites and considerations

Before installing the Tivoli Storage Manager for Space Management client on any system, review the general requirements and considerations.

Installation prerequisites

You must install the Tivoli Storage Manager backup-archive client and you must have authority to install the Tivoli Storage Manager for Space Management client.

- You must have root user authority to install, set up, and use the Tivoli Storage Manager for Space Management client on your workstation.
- You must install and set up the Tivoli Storage Manager backup-archive client before you install the Tivoli Storage Manager for Space Management client. Both the Tivoli Storage Manager for Space Management client and Tivoli Storage Manager 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 Tivoli Storage Manager for Space Management client product is installed. Any files that you place in this directory might be deleted during installation. Do not place the following files into this directory:

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

Installation steps overview

The following steps are an overview of the installation procedure.

- 1. Before installing anything, read all of these general requirements. Read your system installation overview, specific environment requirements, and each installation step to ensure that you are ready to install the Tivoli Storage Manager for Space Management 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 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 Tivoli Storage Manager for Space Management client, stop all activity and do not access any files on file systems to which you added space management. The installation process fails otherwise. Follow the upgrade procedure as described for your operating system. If the steps tell you to uninstall and reinstall the backup archive client and API, you must complete those steps. Uninstalling and reinstalling brings those packages to the same level as the Tivoli Storage Manager for Space Management client.

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

Tivoli Storage Manager for Space Management client GUI requirements and considerations

There are requirements for using the Tivoli Storage Manager for Space Management client graphical user interface (GUI).

- The space management agent and the Tivoli Storage Manager for Space Management client GUI versions must match, otherwise the authentication with the space management agent is not possible.
- The Tivoli Storage Manager node name and password are required to use the Tivoli Storage Manager for Space Management client GUI. The node name and password are the same the Tivoli Storage Manager for Space Management client uses to authenticate with the Tivoli Storage Manager server.
- Before you can connect to a Tivoli Storage Manager client node using the Tivoli Storage Manager for Space Management client GUI, that client node must be registered and authenticated with a Tivoli Storage Manager server.
- The Tivoli Storage Manager for Space Management client GUI requires the following software:

- The Tivoli Storage Manager API, Tivoli Storage Manager backup-archive client, and Tivoli Storage Manager Tivoli Storage Manager for Space Management 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. 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.

Tivoli Storage Manager for Space Management client limitations and considerations

The Tivoli Storage Manager for Space Management client has limitations.

- The Tivoli Storage Manager for Space Management client for AIX GPFS systems is not compatible with the backup-archive client for JFS2. Journal-based backup is not supported when the Tivoli Storage Manager for Space Management client for AIX GPFS systems is installed.
- The Tivoli Storage Manager for Space Management client for GPFS systems manages only file systems belonging to the local (home) GPFS cluster; it does not manage remotely mounted file systems.
- In a GPFS environment, a small file that is less than the size of a GPFS block size can become larger after an HSM migration. GPFS adds meta information to the file during the migration. Because another block on the file system is allocated for the meta information, the space allocated for the file increases. If a file system is filled to its maximum capacity with many small files, it can run out of space during the file migration.
- The Tivoli Storage Manager for Space Management client is not supported on JFS2 Encrypted File Systems (EFS).
- Extended attributes (EAv2) are supported by the Tivoli Storage Manager for Space Management client. However, if you delete a stub file and use the **dsmmigundelete** command to re-create the stub file, the extended attributes are not restored.

Related concepts:

Chapter 3, "Configuring the Tivoli Storage Manager for Space Management client," on page 25

"Space requirement for HSM control files" on page 184

Installation overview of Tivoli Storage Manager for Space Management client for AIX GPFS systems

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

There are several installation limitations for the Tivoli Storage Manager for Space Management client for AIX GPFS systems:

• The Tivoli Storage Manager for Space Management client for AIX GPFS systems is not compatible with the Tivoli Storage Manager for Space Management client

for AIX JFS2 systems or the backup-archive client for JFS2. If you have either of these clients installed and want to install the Tivoli Storage Manager for Space Management client for AIX GPFS systems, you must remove the JFS2 clients.

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

When you install the Tivoli Storage Manager for Space Management 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 **dsmwatchd** command at system startup.
- Removes any statement from the /var/mmfs/etc/gpfsready script file that loads the other space management daemons at GPFS system startup.
- Extracts the HSM modules.
- Adds a statement to the /etc/inittab file that loads the **dsmwatchd** daemon at system startup.
- Adds a statement to the /var/mmfs/etc/gpfsready script file that loads the other space management daemons at GPFS system startup.
- Starts the space management daemons.

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

Table 6. Tivoli Storage Manager for Space Management client for AIX GPFS systems installation packages

Package	Installs	Into this directory
tivoli.tsm.client.ba64.gpfs	The backup-archive client for AIX GPFS	/usr/tivoli/tsm/client/ba/bin
tivoli.tsm.client.hsm.gpfs	The Tivoli Storage Manager for Space Management client for AIX GPFS	/usr/tivoli/tsm/client/hsm/bin
tivoli.tsm.client.api.64bit	The API for AIX	/usr/tivoli/tsm/client/api/bin

For an initial installation, follow these steps:

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

export HSMINSTALLMODE=SCOUTFREE

For information about configuring GPFS integration with the Tivoli Storage Manager for Space Management 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 that are installed with the backup-archive client to configure the Tivoli Storage Manager for Space Management client.
- 3. Install the Tivoli Storage Manager for Space Management client on each HSM node. For AIX clients, see "Installing the Tivoli Storage Manager for Space Management client for AIX systems" on page 12. For Linux clients, see "Installing the Tivoli Storage Manager for Space Management client for Linux x86_64 GPFS systems" on page 18.

- 4. Make sure that after installation, the **dsmrecalld** daemon is running on at least one node.
- 5. Enable the Data Management Application Programming Interface (DMAPI) for GPFS for all file systems to which you plan to add space management. Enable DMAPI only once for each file system.
 - a. Unmount all GPFS file systems on all nodes within the GPFS cluster to which you plan to add space management.
 - b. Activate DMAPI management for the GPFS file systems with the following command: mmchfs *device* -z yes.

For information about GPFS commands, see the *GPFS: Administration and Programming Reference*

For information about GPFS commands, see the IBM Cluster Products information center at http://publib.boulder.ibm.com/infocenter/clresctr/ vxrx/topic/com.ibm.cluster.infocenter.doc/infocenter.html

c. Remount all GPFS file systems on all nodes within the GPFS cluster.

The HSM daemons detect the initial state of each node and assign all nodes an instance number in relation to the GPFS cluster definition.

- 6. On the HSM owner nodes, add space management to each GPFS file system with the **dsmmigfs** command.
- 7. Use the **dsmmigfs enablefailover** command to enable failover of space management on the owner and source cluster nodes that participate in the failover group.

Related tasks:

"Editing the options files by using the backup-archive client" on page 26 **Related reference**:

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

operating systems" on page 44

"dsmmigfs add and update" on page 131

"dsmmigfs sdrreset, enablefailover, and disablefailover" on page 143

Environment requirements for the Tivoli Storage Manager for Space Management client for AIX GPFS systems.

Before you install the Tivoli Storage Manager for Space Management client, your workstation must meet minimum hardware and software requirements.

Hardware requirements

The Tivoli Storage Manager for Space Management client for AIX General Parallel File System (GPFS) systems requires the RS/6000[®] 64-bit pSeries or compatible hardware.

Software requirements

The Tivoli Storage Manager for Space Management client for AIX GPFS systems requires the following software:

- AIX $5L^{M}$ 6.1 PPC in 64-bit kernel mode or AIX 5L 7.1 PPC in 64-bit kernel mode
- GPFS 3.4 or GPFS 3.5

Installing the Tivoli Storage Manager for Space Management client for AIX systems

You can install the Tivoli Storage Manager for Space Management client for AIX systems 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 Tivoli Storage Manager for Space Management client for AIX General Parallel File System (GPFS) systems or the Tivoli Storage Manager for Space Management client for AIX Enhanced Journaled File System (JFS2) systems.

1. Optional: Specify HSMINSTALLMODE=SCOUTFREE.

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

export HSMINSTALLMODE=SCOUTFREE

For information about configuring GPFS integration with the Tivoli Storage Manager for Space Management 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 Tivoli Storage Manager for Space Management 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 21

Upgrading the Tivoli Storage Manager for Space Management client for AIX GPFS systems

You must remove the Tivoli Storage Manager for Space Management GPFS client before installing a new version. You must deactivate the Tivoli Storage Manager for Space Management GPFS client, disable failover, and remove the Tivoli Storage Manager for Space Management GPFS client from all nodes.

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

To upgrade the Tivoli Storage Manager for Space Management GPFS client, complete 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.
- To figure out which node owns which cluster file system, issue the dsmmigfs q -d command on one of the nodes.
- 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 by issuing the command: **dsmmigfs** enablefailover.
- **9**. Take over each file system on its owner node as known from step 4. With distributed HSM, the HSM daemons also run on these nodes, which are designated for the failover feature.

Related concepts:

"Installation overview of Tivoli Storage Manager for Space Management client for AIX GPFS systems" on page 9

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 139

Uninstalling the Tivoli Storage Manager for Space Management client for AIX GPFS systems

Complete this procedure to uninstall the Tivoli Storage Manager for Space Management client for AIX GPFS systems.

Before uninstalling the Tivoli Storage Manager for Space Management client, read the Tivoli Storage Manager for Space Management client for AIX GPFS systems installation overview ("Installation overview of Tivoli Storage Manager for Space Management client for AIX GPFS systems" on page 9).

If you do not remove HSM support from all managed file systems, the data of migrated files are not accessible after you remove the Tivoli Storage Manager for Space Management client.

Follow these steps to uninstall the Tivoli Storage Manager for Space Management client for AIX systems:

- 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:

"Installation overview of Tivoli Storage Manager for Space Management client for AIX GPFS systems" on page 9

Related reference:

"dsmmigfs deactivate, reactivate, and remove" on page 137

Installation overview of the Tivoli Storage Manager for Space Management client for AIX JFS2 systems

Before installing the Tivoli Storage Manager for Space Management client on AIX Enhanced Journaled File System (JFS2) systems, you must review the general requirements, system-specific requirements, and review all installation steps.

There are several installation limitations for the Tivoli Storage Manager for Space Management client for AIX JFS2 systems:

- The Tivoli Storage Manager for Space Management client for AIX JFS2 systems is not compatible with the Tivoli Storage Manager for Space Management client for AIX General Parallel File System (GPFS) systems or the backup-archive client for GPFS. If you have either of these clients installed and want to install the Tivoli Storage Manager for Space Management client for AIX JFS2 systems, you must remove the GPFS clients.
- On AIX 6.1 and 7.1, the Tivoli Storage Manager for Space Management client can be installed in the global partition and supports transparent recall for both global and local workstation partitions (WPARs). Using HSM commands from a local WPAR is not supported. You cannot install the Tivoli Storage Manager for Space Management client in a local WPAR.

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

Table 7. Tivoli Storage Manager for Space Management client for AIX JFS2 systems 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 Tivoli Storage Manager for Space Management 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 Tivoli Storage Manager for Space Management client. American English messages are already included in the Tivoli Storage Manager for Space Management client.	/usr/tivoli/tsm/client/lang/bin

Related reference:

"dsmmigfs add and update" on page 131

"Adding space management to workload partitions on AIX V6.1 and V7.1 operating systems" on page 44

Environment requirements for the Tivoli Storage Manager for Space Management client for AIX JFS2 systems

Before installing the Tivoli Storage Manager for Space Management client on an AIX JFS2 system, you must review the hardware and software requirements.

Hardware requirements

The Tivoli Storage Manager for Space Management client for AIX JFS2 systems requires the following hardware:

• RS/6000 64-bit pSeries, or compatible hardware

Software requirements

The Tivoli Storage Manager for Space Management client for AIX JFS2 systems 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 Tivoli Storage Manager for Space Management client for AIX systems

You can install the Tivoli Storage Manager for Space Management client for AIX systems 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 Tivoli Storage Manager for Space Management client for AIX General Parallel File System (GPFS) systems or the Tivoli Storage Manager for Space Management client for AIX Enhanced Journaled File System (JFS2) systems.

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

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

export HSMINSTALLMODE=SCOUTFREE

For information about configuring GPFS integration with the Tivoli Storage Manager for Space Management 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 Tivoli Storage Manager for Space Management 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 21

Installation overview for the Tivoli Storage Manager for Space Management client for Linux x86_64 GPFS systems

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

Note:

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

When you install the Tivoli Storage Manager for Space Management 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 lists the packages available on the installation media:

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 Tivoli Storage Manager for Space Management client for Linux x86_64 (64-bit only)	/opt/tivoli/tsm/client/hsm/bin
TIVsm-msg. <i>lang</i> .x86_64.rpm	The message package for language (where <i>lang</i> is the language code)	/opt/tivoli/tsm/client/lang

Table 8. Linux GPFS available packages

For an initial installation, follow these steps:

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

export HSMINSTALLMODE=SCOUTFREE

For information about configuring GPFS integration with the Tivoli Storage Manager for Space Management 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 that are installed with the backup-archive client to configure the Tivoli Storage Manager for Space Management client.
- 3. Install the Tivoli Storage Manager for Space Management client on each HSM node. For AIX clients, see "Installing the Tivoli Storage Manager for Space Management client for AIX systems" on page 12. For Linux clients, see "Installing the Tivoli Storage Manager for Space Management client for Linux x86_64 GPFS systems" on page 18.
- 4. Make sure that after installation, the **dsmrecalld** daemon is running on at least one node.
- 5. Enable the Data Management Application Programming Interface (DMAPI) for GPFS for all file systems to which you plan to add space management. Enable DMAPI only once for each file system.
 - a. Unmount all GPFS file systems on all nodes within the GPFS cluster to which you plan to add space management.
 - b. Activate DMAPI management for the GPFS file systems with the following command: mmchfs *device* -z yes.

For information about GPFS commands, see the *GPFS: Administration and Programming Reference*

For information about GPFS commands, see the IBM Cluster Products information center at http://publib.boulder.ibm.com/infocenter/clresctr/vxrx/topic/com.ibm.cluster.infocenter.doc/infocenter.html

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

Related tasks:

"Editing the options files by using the backup-archive client" on page 26

Related reference:

"dsmmigfs add and update" on page 131

"dsmmigfs sdrreset, enablefailover, and disablefailover" on page 143

Environment requirements for the Tivoli Storage Manager for Space Management client for Linux x86_64 GPFS systems

Before you install the Tivoli Storage Manager for Space Management client on a Linux x86_64 General Parallel File System (GPFS) system, your workstation must meet minimum communication, hardware, and software requirements.

Software requirements

The Tivoli Storage Manager for Space Management 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

Installing the Tivoli Storage Manager for Space Management client for Linux x86_64 GPFS systems

You can install the Tivoli Storage Manager for Space Management client on Linux x86_64 General Parallel File System (GPFS) systems using the product DVD.

Follow these steps:

- 1. Log in as the root user and mount the DVD to the /cdrom directory if applicable.
- 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 **dsmmonitord** and **dsmscoutd** automatic migration daemons. Disabling these daemons conserves system resources. To disable the automatic migration daemons, start this command in a shell:

export HSMINSTALLMODE=SCOUTFREE

For information about configuring GPFS integration with the Tivoli Storage Manager for Space Management 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.

To install	Issue this command
API	Install the API files in the /opt/tivoli/tsm/client/api/bin64 directory.
	rpm -i TIVsm-API64.x86_64.rpm
Backup-archive client (CLI and GUI) web	Install the backup-archive client (command line, GUI, and web client) in the /opt/tivoli/tsm/client/ba/bin directory.
client Administrative client	rpm -i TIVsm-BA.x86_64.rpm
	To circumvent the dependence check, you can use thenodeps option, but then you must check the dependencies manually.
	To install the backup-archive client (command line, GUI, and web client) and administrative client in the /opt/tivoli/tsm/client/ba/ bin directory:
	rpm -i TIVsm-BA.x86_64.rpm
Tivoli Storage Manager for Space Management client	To install the following Tivoli Storage Manager for Space Management client files in the /opt/tivoli/tsm/client/hsm/bin directory.
	Tivoli Storage Manager for Space Management client commands
	Space management agent (hsmagent)
	• IBM Tivoli Enterprise Space Management Console (Tivoli Storage Manager for Space Management client GUI)
	rpm -i TIVsm-HSM.x86_64.rpm

Table 9. Package names and descriptions

Related reference:

"hsmdisableautomigdaemons" on page 99

Upgrading the Tivoli Storage Manager for Space Management client for Linux x86_64 GPFS systems

You must remove the Tivoli Storage Manager for Space Management GPFS client before installing a new version. You must deactivate the Tivoli Storage Manager for Space Management GPFS client, disable failover, and remove the Tivoli Storage Manager for Space Management GPFS client from all nodes.

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

To upgrade the Tivoli Storage Manager for Space Management GPFS client, complete 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.
- 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 by issuing the command: dsmmigfs enablefailover.
- **9**. Take over each file system on its owner node as known from step 4. With distributed HSM, the HSM daemons also run on these nodes, which are designated for the failover feature.

Related concepts:

"Installation overview for the Tivoli Storage Manager for Space Management client for Linux x86_64 GPFS systems" on page 16

Related reference:

```
"dsmmigfs globaldeactivate and globalreactivate" on page 138
```

```
"dsmmigfs sdrreset, enablefailover, and disablefailover" on page 143
```

```
"dsmmigfs query" on page 139
```

"dsmmigfs deactivate, reactivate, and remove" on page 137

Uninstalling the Tivoli Storage Manager for Space Management client for Linux x86_64 GPFS systems

You must follow a specific procedure to uninstall the Tivoli Storage Manager for Space Management client for Linux x86_64 GPFS systems.

To uninstall the Tivoli Storage Manager for Space Management client for Linux x86_64 GPFS systems, complete the following steps.

- 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:

"Installation overview for the Tivoli Storage Manager for Space Management client for Linux x86_64 GPFS systems" on page 16

Related reference:

"dsmmigfs deactivate, reactivate, and remove" on page 137

Registering your workstation with a Tivoli Storage Manager server

After installing the Tivoli Storage Manager for Space Management 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 Tivoli Storage Manager for Space Management client 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 provided 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 Tivoli Storage Manager for Space Management 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 Tivoli Storage Manager for Space Management 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 164

The Tivoli Storage Manager for Space Management client GUI

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

The Tivoli Storage Manager for Space Management client GUI can be used on any HSM system where the HSM agent is running.

The Tivoli Storage Manager for Space Management client GUI can also be used on Windows systems so you can remotely administer more than one Tivoli Storage Manager for Space Management client from a single point. Download the Windows package from ftp://public.dhe.ibm.com/storage/tivoli-storage-management/.

Before using Tivoli Storage Manager for Space Management client GUI, check the following prerequisites:

- 1. Java Runtime Environment (JRE) Version 6 is installed.
- 2. The Tivoli Storage Manager for Space Management client is installed and the **hsmagent** is started on any computer that is managed through the Tivoli Storage Manager for Space Management client GUI.
- **3**. The Tivoli Storage Manager client node used for 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 Tivoli Storage Manager for Space Management client GUI, access its online help from the **Help** menu.

Verifying that the HSM agent is running

In order to use the Tivoli Storage Manager for Space Management client GUI to manage an HSM system, the HSM agent must be running.

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

- Log in or telnet to the HSM system that you want to manage with the Tivoli Storage Manager for Space Management client GUI by issuing the command: telnet HSM_machine_name
- 2. Verify that the Tivoli Storage Manager for Space Management client version installed matches the version of the Tivoli Storage Manager for Space Management client 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.

```
dsmadmc 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 complete 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 Tivoli Storage Manager for Space Management client GUI

When you use the Tivoli Storage Manager for Space Management client 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 Tivoli Storage Manager for Space Management client GUI:

- On Windows, click Start and select Programs > Tivoli Storage Manager > Space Management Console. On AIX or Linux systems, run the dsmsmj command where the Tivoli Storage Manager for Space Management client is supported and installed. The Welcome Page opens.
- 2. 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 computer. You must authenticate with at least one client node to retrieve the saved password.
- **3**. 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 the appropriate 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 for any of the following tasks:
 - Add Space Management
 - Remove Space Management
 - Activate Space Management
 - Deactivate Space Management
 - Modify specific file system properties
- 6. Click the appropriate button to display a graphical overview of the file systems on the selected client node.

Chapter 3. Configuring the Tivoli Storage Manager for Space Management client

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

The Tivoli Storage Manager for Space Management client is configured by setting options in the dsm.sys and dsm.opt files. The options generally affect all space-managed file systems, and determine the following things:

- The 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 the copy expires.
- Whether to use two-way orphan check reconciliation (for GPFS only)

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

- On AIX: /usr/tivoli/tsm/client/ba/bin64
- On Linux x86_64: /opt/tivoli/tsm/client/ba/bin

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

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

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 Tivoli Storage Manager for Space Management client shares the following common files and code with the AIX and Linux backup-archive clients:

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

• Server file spaces

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

For information about configuring GPFS integration with the Tivoli Storage Manager for Space Management 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, "Options files reference," on page 95

Displaying Tivoli Storage Manager for Space Management client options

You can display information about Tivoli Storage Manager for Space Management client options either from the 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 Tivoli Storage Manager for Space Management client GUI, complete the following steps.

Note: You can use the Tivoli Storage Manager for Space Management client GUI's online help to get more detailed information. Select **Help Topics** from the **Help** menu, by pressing F1, or by clicking the **?** icon in the upper right of the Tivoli Storage Manager for Space Management client 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 147

Editing the options files by using the backup-archive client

The Tivoli Storage Manager for Space Management client shares the option files, dsm.opt and dsm.sys, with the IBM Tivoli Storage Manager for AIX and Linux backup-archive clients. To edit the options files, you can use the Preferences Editor window of the backup-archive client GUI. You can also edit the files in a text editor.

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 that you define in the dsm.sys file for the Tivoli Storage Manager for Space Management 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
- 4. Options that you enter locally in your options file
- **5**. Options that are received from the server with client options that the server does not enforce (the client can override the value)
- 6. Default option values

For options that you can set in the dsm.sys file, see Chapter 11, "Options files reference," on page 95

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

- 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 **commethod** entry to specify the communication method to use for client and server communications.
- 4. Issue a value for each option and remove the leading asterisk (*). You can specify options for more than one server.

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

1500 server3.almaden.ibm.com generate
/usr/bin/xsend root
system tsm
steiner chron wang nguyen
/adm/tsm/backup.excl
/bin/cat
SNAlu6.2
raptor
appcdel
appc
generate
/usr/bin/xsend root
system tsm
sullivan tang stewart
/adm/tsm/migrate.excl
/bin/cat

Editing the options file dsm.opt

Most of the options in the dsm.opt file affect the backup-archive client. However, some options affect the Tivoli Storage Manager for Space Management client.

Table 10 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 the Tivoli Storage Manager for Space Management client. The Tivoli Storage Manager for Space Management client operates only on local file systems.

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

Table 10. Tivoli Storage Manager for Space Management client options in the dsm.opt file

Option	Default	Description
defaultserver	The server that you identify in the first stanza of the dsm.sys file.	Specifies the name of the default Tivoli Storage Manager server to contact for space management services when you do not specify a server name on the migrateserver option. The value of defaultserver in the dsm.sys file overrides defaultserver in the dsm.opt file. See " defaultserver " on page 98.
hsmgroupedmigrate	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 txngroupmax option, which is set on the Tivoli Storage Manager server. See "hsmgroupedmigrate" on page 102.
migrateserver	The value of defaultserver option.	Specifies the Tivoli Storage Manager server to which you want to migrate files from your client node. Specify only 1 server for each client node. The value of migrateserver in the dsm.sys file overrides migrateserver in the dsm.opt file. See " migrateserver " on page 111.

Table 10. Tivoli Storage Manager for Space Management client options in the dsm.opt file (continued)

Option	Default	Description
restoremigstate	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 " restoremigstate " on page 113.

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

Related tasks:

"Setting environment variables" on page 34

Configuring the Tivoli Storage Manager for Space Management client to connect to a secondary Tivoli Storage Manager server

If the primary Tivoli Storage Manager server for the Tivoli Storage Manager for Space Management client is unavailable, you can manually configure the Tivoli Storage Manager for Space Management client to connect to a secondary server. You can recall files from the secondary Tivoli Storage Manager server but cannot migrate files to the secondary server.

The primary Tivoli Storage Manager server for the Tivoli Storage Manager for Space Management client must be one that replicates client node data.

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

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

The Tivoli Storage Manager for Space Management client, however, does not automatically fail over to the secondary server. You must manually edit the dsm.sys file to connect to the secondary server. Any secondary server information in the **replservername** stanza and **myreplicationserver** option is ignored by the Tivoli Storage Manager for Space Management client.

To configure the Tivoli Storage Manager for Space Management client to connect to the secondary server, complete the following steps:

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

Servername	lifeboat server	
COMMmethod	-	TCPip
TCPPort		1500
TCPServerado	lress	<pre>server4.almaden.ibm.com</pre>
Passwordacce	ess	generate
Mailprog		/usr/bin/xsend root
Groups		system tsm
Users		steiner chron wang nguyen
Inclexcl		/adm/tsm/backup.excl
ERRORProg		/bin/cat

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

DEFAULTServer	lifeboat server
MIGRATEServer	lifeboat_server

- 3. For non-root users, edit the dsm.opt file and update for value in the defaultserver option to point to the secondary server stanza that you added in the dsm.sys file.
 DEFAULTServer lifeboat server
- 4. Restart the Tivoli Storage Manager for Space Management client, including all space management daemons.

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

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

To switch back to the primary Tivoli Storage Manager server, repeat the steps, but change the value of the **servername** option from the target server to the source server.

Related tasks:

"Stopping the space management daemons" on page 91

Optional setup features

You can specify which files are eligible for migration. You can use clustered file systems and LAN-free data transfer, start the Tivoli Storage Manager for Space Management client graphical user interface, and set logging options.

Assignment of management classes to files

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

The 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 with one or more include statements in your include-exclude options file.

A management class can contain a backup copy group and an archive copy group. Copy groups contain attributes that control the generation, destination, and expiration of backup versions of files and archived copies of files. 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 11 on page 31 lists the space management attributes and their defaults that might be included in a management class.

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

Table 11. Tivoli Storage Manager for Space Management client attributes in a management class

Table 12 identifies appropriate management classes for some tasks.

Table 12. Assigning management classes to your files

Task	Assignment
You want to migrate and back up a file.	Assign a management class to a file with space management attributes and backup-archive copy groups that you want to use for that file. Assign only one management class to a specific file.
Your client node communicates with the same server for both space management and backup-archive services.	Assign a management class that contains space management attributes and backup-archive copy groups that you assigned to the backup-archive client to use for that file.
You migrate files to one server and back up and archive files to one or more different servers:	 Specify a different include-exclude options file for each server. Assign only one management class to a file in a specific include-exclude options file. Assign different management classes to files in different include-exclude options files. 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. 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 mgmt2b for the same file. That management class must contain appropriate space management

Displaying management class information

You can display management class information from the Tivoli Storage Manager for Space Management client GUI for the default migration server only. If your client node contacts one or more extra 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 Tivoli Storage Manager for Space Management client GUI to display policy and management classes that you can assign to your files, follow these steps:

Note: To access the Tivoli Storage Manager for Space Management client GUI online help, select the **Help Topics** from the **Help** menu, press F1, or click the **?** icon.

- 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 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 147

Incude-exclude file options

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

Note:

- If you do not create an include-exclude options file, all files are considered for backup services and the default management class is used.
- On AIX GPFS and Linux x86_64 GPFS file systems do not use the **inclexcl** option in a failover environment. Unlike dsm.opt and dsm.sys files, include-exclude options files are not shared between different nodes of a failover group. Add the include-exclude list directly to the dsm.sys file or verify that the include-exclude list files match on all nodes that participate in the local failover group or node set.

For all other file systems, use the **inclexcl** 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 **inclexcl** option in each stanza. Each **inclexcl** option can point to a different include-exclude options file. The files that you create must reside in a directory to which all users on your workstation have read access.

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

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

The following options include or exclude objects from space management.

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

Table 13. Include and Exclude Statements

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

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

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

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

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

Creating an include-exclude list

The Tivoli Storage Manager for Space Management client shares the include-exclude list with the Tivoli Storage Manager for AIX and Linux backup-archive clients. You can use the backup-archive GUI or the command line to create and edit your include-exclude list.

Note:

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

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

- 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, complete the following steps.

- 1. Create an empty file in any directory to which all users on your workstation have read access.
- 2. Enter your include and exclude statements.
- 3. In the dsm.sys file, use the **inclexcl** 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.

The following environment variables identify Tivoli Storage Manager for Space Management client files.

Table 14. Environment Variables

Variable	Description
DSM_DIR	Points to the resource files, the dsm.sys file, and the executable file, dsmtca. If you do not set DSM_DIR , the Tivoli Storage Manager for Space Management client searches for the executable files in the installation directory.

Table 14. Environment Variables (continued)

Variable	Description
DSM_CONFIG	Points to your dsm.opt file.
	 If you do not set DSM_CONFIG, the Tivoli Storage Manager for Space Management client searches for the options file in the directory to which DSM_DIR points.
	 If you do not set DSM_DIR, the Tivoli Storage Manager for Space Management client searches for the options file in the installation directory.
DSM_LOG	Points to the directory where you want the dsmerror.log file to reside. The error log file contains information about any errors that occur during processing. This log file helps Tivoli Customer Service diagnose severe errors.
	 If you define DSM_DIR but you do not define DSM_LOG, messages are written to dsmerror.log in the directory that you specified in DSM_DIR.
	 If you do not define DSM_LOG or DSM_DIR, error messages are written to dsmerror.log in the current directory. You receive a warning message if Tivoli Storage Manager cannot write messages to the log file. Processing continues.

Setting Bourne and Korn shell variables

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

For example:

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

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

Setting C shell variables

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

For example:

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

The /home/hsmuser/dsm.opt path identifies the path and file name for your dsm.opt file. The /home/hsmuser path identifies the directory where you want to store the dsmerror.log file.

LAN-free data transfer for HSM

The Tivoli Storage Manager for Space Management client supports LAN-free data transfer, which shifts the movement of client data from the communications network to a storage area network (SAN). Shifting the client data movement from the communications network to a SAN decreases the load on the 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 you establish LAN-free support in your HSM environment, you must check the prerequisites.

To enable LAN-free support the following prerequisites are 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 you install and configure the Tivoli Storage Manager Managed System for SAN feature on the client workstation, use options to enable LAN-Free data transfer.

Table 15 lists options that 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*.

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

Table 15. LAN-free data transfer options

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

To include files:

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

To exclude files:

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

HSM 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 Tivoli Storage Manager for Space Management 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. Across different nodes, the server stanza must have the following characteristics:

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

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

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

To enable distributed migration and recall capabilities follow these steps:

- 1. Install the Tivoli Storage Manager for Space Management client on all participating GPFS nodes in the cluster environment.
- 2. Register each Tivoli Storage Manager for Space Management 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 do not require HSM. Nodes that are used only to access data of space-managed file systems are not required to be registered on the 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.

Configuration examples are available in the Tivoli Storage Manager Space Management-GPFS integration field guide.

For information about configuring GPFS integration with the Tivoli Storage Manager for Space Management 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:

"Installation overview for the Tivoli Storage Manager for Space Management client for Linux x86_64 GPFS systems" on page 16

"Installation overview of Tivoli Storage Manager for Space Management client for AIX GPFS systems" on page 9

Related reference:

"hsmdistributedrecall" on page 100

Limitations of HSM support for AIX GPFS and Linux x86_64 GPFS

Be aware of the following limitations of HSM support for AIX General Parallel Files System (GPFS) and Linux x86_64 GPFS file systems.

GPFS support is limited in the following ways:

- 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 integrated with the backup-archive client support. For example, the Tivoli Storage Manager for Space Management client refers to the *file_system*/.SpaceMan/hsmfsconfig.xml file to determine which server to use for a file system. The Tivoli Storage Manager for Space Management client might contact a different server for each file system. In contrast, the backup-archive client uses other methods to determine which server to use for a backup process.

HSM space management agent configuration

The HSM agent program, hsmagent (also called the space management agent), is the communication endpoint for the Tivoli Storage Manager for Space Management client GUI. After you start the HSM agent, the Tivoli Storage Manager for Space Management client GUI can connect to it on the port that is 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 Tivoli Storage Manager for Space Management client to connect the Tivoli Storage Manager for Space Management client GUI with the HSM agent.

If you 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 in /usr/tivoli/tsm/client/hsm/bin. On Linux x86_64 GPFS, the file is 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 1 (Table 16 on page 39 describes the options):

Figure 1. 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 16. HSM agent options

Parameter	Description
PortNumber	The PortNumber option specifies on which port the Tivoli Storage Manager for Space Management client GUI can connect to the space management agent. Valid values are 1500 - 32152.
AgentTimeOut	Do not change the AgentTimeOut option except on the advice of an IBM service representative.
SessionTimeOut	Do not change the SessionTimeOut option except on the advice of an IBM service representative.
GuiOptionsFile	Do not change the GuiOptionsFile option except on the advice of an IBM service representative.
Tracing	Do not change the Tracing, TraceFile, and TraceFlags options except on the advice of an IBM service representative.

Logs for HSM activity and error messages

The Tivoli Storage Manager for Space Management client logs HSM activity and error messages. An HSM log contains information about file migration and recall, threshold migration, reconciliation, and starting and stopping the HSM daemon. An error log records error messages.

Logs for HSM activity

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

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

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

- hsmlogeventflags
- hsmlogmax
- hsmlogname
- hsmlogretention
- hsmlogsampleinterval

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

- dsmmigrate
- dsmrecall
- dsmmigundelete
- dsmmigfs
- dsmdf

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

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

Logs for error messages

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

- errorlogmax
- errorlogname
- errorlogretention

You can briefly change the error log file name if you specify the **errorlogname** parameter on any Tivoli Storage Manager for Space Management client command. The new log captures the errors of only the operation that was started by the command.

Related reference:

"hsmlogeventflags" on page 103

- "hsmlogmax" on page 104
- "hsmlogname" on page 105

"hsmlogretention" on page 105

"hsmlogsampleinterval" on page 106

"errorlogname" on page 98

"dsmmigrate" on page 149

"dsmrecall" on page 156

"dsmmigundelete" on page 152

"dsmdf" on page 125

"dsmmigfs query" on page 139

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 Tivoli Storage Manager for Space Management client can work only 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 Tivoli Storage Manager for Space Management client is adding space management to your file systems, do not attempt to access any files in your file systems. Do not perform any tasks against your file systems.

When you add space management to a file system, the Tivoli Storage Manager for Space Management client completes the following tasks:

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

Remember: The .SpaceMan directory is not processed by the 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 **dsmwatchd** daemon starts at system startup with an entry in the etc/inittab file or with the **initctl** service. Ensure that DMAPI is enabled on all GPFS file systems that the Tivoli Storage Manager for Space Management client manages. Issue the following command to query this information: /usr/lpp/mmfs/bin/mmlsfs *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 the value for the DMAPI enablement to YES only if the file system is unmounted on all nodes of the cluster. When DMAPI is enabled, the file system can be mounted only if a dsmrecalld daemon is set up on one of the cluster nodes within the GPFS cluster.

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 Tivoli Storage Manager for Space Management client on several AIX or Linux x86_64 cluster nodes are candidates for managing a GPFS file system, the preferred node has precedence. If failover occurs, the Tivoli Storage Manager for Space Management client that manages the file system is not the Tivoli Storage Manager for Space Management client on the preferred node.

For information about configuring GPFS integration with the Tivoli Storage Manager for Space Management 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.

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

- Move stub files as native stub files (without the migrated data) to other file systems, even if the other file system is space-managed.
- Use the Tivoli Storage Manager raw device backup utility or any other image-based backup utility to restore a space-managed file system to a system other than the original system.
- Use space-managed file systems within cluster replication tools.
- Use any other tools that transfer images between systems.

The following are more considerations:

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

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

Related tasks:

"Adding space management to nested file systems" on page 43

"Adding space management to an exported file system" on page 44 **Related reference**:

Appendix A, "Control files in the .SpaceMan directory," on page 183

Adding space management to file systems

You can add space management by issuing a command or by using the Tivoli Storage Manager for Space Management client GUI.

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 Tivoli Storage Manager for Space Management 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.

If you do not use the GPFS policy engine to control migration, specify the space thresholds with the **dsmmigfs add** command or the **dsmmigfs update** command or with the Tivoli Storage Manager for Space Management client GUI.

To add space management to your local file systems from the command line, use the **dsmmigfs add** command.

To add space management to a file system with the Tivoli Storage Manager for Space Management client GUI, follow these steps:

- 1. From the Tivoli Storage Manager for Space Management client GUI 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 Tivoli Storage Manager for Space Management client displays a list of file systems next to the **Client Nodes** tree.
- 2. Select a file system with status Not Managed, and click **Manage**. The status of the file system becomes Active.
- **3**. If you want to change the default space management settings, click **File System Properties**. Make and save any needed changes.

Related reference:

"dsmmigfs add and update" on page 131

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.

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

Related tasks:

"Adding space management to file systems" on page 42

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 must 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 jfs2.
- 2. Select Change/Show Characteristics of an Enhanced 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 Tivoli Storage Manager for Space Management 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 supports only the NFS (Network File System) interface to export a file system.

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.
- 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 131

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 Tivoli Storage Manager for Space Management 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 Tivoli Storage Manager for Space Management client on the node.
- 2. Use the dsmmigfs takeover command to regain control of the file systems.

Related reference:

"dsmmigfs takeover" on page 145

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

AIX V6.1 and V7.1 workload partitions (WPARs) act and look like stand-alone systems and provide an isolated environment for enterprise applications and data. A Tivoli Storage Manager for Space Management 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.*

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

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

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

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

The Tivoli Storage Manager for Space Management 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 update** command.

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

For information about configuring GPFS integration with the Tivoli Storage Manager for Space Management 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.

If you do not use the GPFS policy engine to control migration, specify the settings with the **dsmmigfs update** 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 131

Updating settings from the Tivoli Storage Manager for Space Management client GUI

You can update space management settings to change the way the Tivoli Storage Manager for Space Management 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 Tivoli Storage Manager for Space Management client maintains on your local file system
- The minimum size (in megabytes) for a file to qualify for partial file recall
- The minimum size (in megabytes) for a file to qualify for streaming recall mode
- The maximum number of files in the space-managed file system.

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

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

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

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

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

- If you change the high and low thresholds or the premigration percentages for a file system, the new values take effect immediately.
- If you change the stub file sizes, the new values take effect only for files that are migrated after you change the settings.

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

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

For information about configuring GPFS integration with the Tivoli Storage Manager for Space Management 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 Tivoli Storage Manager for Space Management client," on page 25

Related reference: "dsmmigfs query" on page 139

"dsmmigfs add and update" on page 131

"dsmmigrate" on page 149

Minimum migration file size

The Tivoli Storage Manager for Space Management client does not migrate a file unless the migration 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.

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

Related reference:

"dsmmigfs add and update" on page 131

"minmigfilesize" on page 112

Migration threshold percentages

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

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

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

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

If the high threshold and the low threshold are the same, space usage must exceed the low threshold before threshold migration begins. When setting the realistic-minimum low-threshold option, remember that the .SpaceMan directory (which contains system-specific control files created by 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 Tivoli Storage Manager for Space Management 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 131

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 Tivoli Storage Manager for Space Management 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 Tivoli Storage Manager for Space Management 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 Tivoli Storage Manager for Space Management 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 Tivoli Storage Manager for Space Management 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 Tivoli Storage Manager for Space Management 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 Tivoli Storage Manager for Space Management 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 131

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 999999999999999.

- 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 99999999999999, 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 Tivoli Storage Manager for Space Management client migrates and premigrates your files from that file system until the total number of MB that migrate and premigrate equals 20 GB.

Check with your 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 131

Stub file size

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

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

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

For Tivoli Storage Manager for Space Management clients on AIX Enhanced Journaled Files Systems (JFS2), you can specify a multiple of the file system fragment size. The default value is the file system fragment size.

For Tivoli Storage Manager for Space Management clients on AIX GPFS and Linux $x86_{64}$ GPFS file systems, you can specify 0 or a multiple of the file system block size. The default value is 0.

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

"dsmmigfs add and update" on page 131

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 99999999. A value of 0 disables the asynchronous option and is the default.

Related reference:

"dsmmigfs add and update" on page 131

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 4

"File migration eligibility" on page 62

Maximum number of files

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

You can specify a maximum number of files value from 0 - 99999999999999999.

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

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

For information about configuring GPFS integration with the Tivoli Storage Manager for Space Management 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:

"Streaming recall mode" on page 79

Related reference:

- "dsmmigfs add and update" on page 131
- "dsmmigfs query" on page 139

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 Tivoli Storage Manager for Space Management client distributes files to all servers in a round-robin fashion. You can customize the distribution scheme to suit your environment.

A single Tivoli Storage Manager for Space Management client can use only HSM daemons (**dsmmonitord** and **dsmscoutd**) or GPFS policy to manage automatic migration. A Tivoli Storage Manager for Space Management client cannot use both HSM daemons and GPFS policy to manage automatic migration.

Related tasks:

"Restoring a file system in an environment that is managed by multiple Tivoli Storage Manager servers" on page 75

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

Enable multiple-server support for each Tivoli Storage Manager for Space Management client in the GPFS node, 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 a Tivoli Storage Manager for Space Management 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 Tivoli Storage Manager for Space Management client can use only HSM daemons (**dsmmonitord** and **dsmscoutd**) or GPFS policy to manage automatic migration. A Tivoli Storage Manager for Space Management client cannot use both HSM daemons and GPFS policy to manage automatic migration.

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

For information about configuring GPFS integration with the Tivoli Storage Manager for Space Management 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 Products information center: http://publib.boulder.ibm.com/infocenter/clresctr/vxrx/topic/ com.ibm.cluster.gpfs.doc/gpfs_faqs/gpfsclustersfaq.html.

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

- Set option hsmdisableautomigdaemons=YES for each Tivoli Storage Manager for Space Management clientin the GPFS cluster. Setting hsmdisableautomigdaemons=YES disables the automatic migration daemons dsmmonitord and dsmscoutd. The dsmmonitord and dsmscoutd daemons do not drive automatic migration in a multiple-server environment. If you do not disable the dsmmonitord and dsmscoutd daemons, they continue to run and use system resources.
- 2. Set option hsmmultiserver=YES for each Tivoli Storage Manager for Space Management clientin 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_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 **dsmreconcile** command to run a special reconciliation that couples the files with the server.

Optional: Add additional servers to manage a file system. For each additional server, issue the following command: dsmmigfs addmultiserver -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 135

"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.

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

The Tivoli Storage Manager for Space Management 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 131

"dsmmigfs addmultiserver, querymultiserver, and removemultiserver" on page 135

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:

- 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 you run the command, files are not backed up or migrated to the server that was removed. You can still recall migrated copies and restore backup copies from the server that was removed.
- 2. Run the dsmMultiServerRemove.pl script.

The dsmMultiServerRemove.pl script does the following three tasks:

a. Recalls all migrated files that are coupled with the Tivoli Storage Manager server that was removed to the specified file system.

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

b. Reconciles the specified file system with the Tivoli Storage Manager server that was removed. The following reconciliation tasks are completed:
- The Tivoli Storage Manager for Space Management 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 Tivoli Storage Manager for Space Management 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 commands:

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 command: **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 135

"dsmRemoveServer.pl" on page 171

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 the primary Tivoli Storage Manager server for the Tivoli Storage Manager for Space Management client is unavailable, you can manually configure the Tivoli Storage Manager for Space Management client to connect to a secondary server. You can recall files from the secondary Tivoli Storage Manager server but cannot migrate files to the secondary server.

The primary Tivoli Storage Manager server for the Tivoli Storage Manager for Space Management client must be one that replicates client node data.

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

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

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

The Tivoli Storage Manager for Space Management client, however, does not automatically fail over to the secondary server. You must manually edit the dsm.sys file to connect to the secondary server. Any secondary server information in the **replservername** stanza and **myreplicationserver** option is ignored by the Tivoli Storage Manager for Space Management client.

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

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

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

To switch back to the primary Tivoli Storage Manager server, repeat the steps, but change the value of the **servername** option from the target server to the source server.

Deactivating space management

Use a command or the Tivoli Storage Manager for Space Management client 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 complete 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 while importing and exporting to another server.

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 Tivoli Storage Manager for Space Management client 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 with the command line or the Tivoli Storage Manager for Space Management client 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 with the Tivoli Storage Manager for Space Management client 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. A client node can be in one of the following states:

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

You can also complete 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.

- 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 **Reactivate**.

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 with a single command or with the Tivoli Storage Manager for Space Management client 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 is stopped, and that the file system is not being accessed.

When you remove space management from your file system, the Tivoli Storage Manager for Space Management client does the following tasks:

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

- Deletes the .SpaceMan directory from your file system.
- Updates information for your native file system.

To remove space management from an active file system with the Tivoli Storage Manager for Space Management client GUI, follow these steps:

- 1. From the Manage Resourceswindow, 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.
- Select an Active file system from the table in the work area, and click Unmanage. 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 complete 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

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

Migration types

The Tivoli Storage Manager for Space Management client provides automatic and selective migration. Automatic migration is triggered by space-usage thresholds (threshold migration) or an out-of-space condition (demand migration).

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

Threshold migration

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

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

For information about configuring GPFS integration with the Tivoli Storage Manager for Space Management 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.

Demand migration

Demand migration responds to an out-of-space condition on your local file system. Demand migration starts automatically if your file system runs out of space. For HSM on AIX JFS2, as files are migrated, space becomes available on your file system. 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.

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

File migration eligibility

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

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

- The file is a regular file that you previously did not migrate. Character special files, block special files, FIFO special files (named pipe files), or directories are not migrated.
- The file is a resident or premigrated file on a file system for which space management is active.
- The file is not excluded from space management in your include-exclude options file.
- The file size is equal to or greater than the value of the **minmigfilesize** option.

Restriction: This criterion does not apply to premigration.

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

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

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

The Tivoli Storage Manager backup-archive client uses the **preservelastaccessdate** option to determine whether to reset the last access date of any specified files after a backup or archive operation. By default, the backup-archive client does not reset the last access dates to their original values after the backup or archive operation.

Related concepts:

Chapter 1, "Tivoli Storage Manager for Space Management client overview," on page 1

"Migrating files overview" on page 3

Related reference:

"dsmdf" on page 125

"dsmdu" on page 126

"hsmenableimmediatemigrate" on page 101

File premigration

For faster migration, the Tivoli Storage Manager for Space Management 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. The next time you need free space on your local file system, the Tivoli Storage Manager for Space Management client just changes the status of premigrated files to migrated files. No additional time is required to copy the files to storage.

The Tivoli Storage Manager for Space Management 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 Tivoli Storage Manager for Space Management client premigrates files each time it completes automatic migration if the following conditions are true:

- The file system contains additional files that are eligible for automatic migration.
- The premigration percentage that you set for your file system is not 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 high threshold percentage and the low threshold percentage. You can change the premigration percentage at any time.

You can manually premigrate files with the **dsmmigrate** command and the **premigrate** option.

Related reference:

"hsmgroupedmigrate" on page 102

"dsmmigrate" on page 149

Automatic file migration

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

The Tivoli Storage Manager for Space Management client provides two types of automatic migration: threshold and demand.

Threshold migration maintains a specific level of free space on your local file system. The space monitor daemon checks space usage on your local file systems at intervals that you specify. When space usage reaches the high threshold that you set for a file system, migration automatically sends eligible files to storage. When space usage reaches the low threshold that you set for a file system, migration stops. For example, assume that you set the high threshold for your file system to 80 percent and the low threshold to 70 percent. Files begin migrating to 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 **pmpercentage** parameter of the **dsmmigfs** command.

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

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

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

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

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

For information about configuring GPFS integration with the Tivoli Storage Manager for Space Management 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 131

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 Tivoli Storage Manager for Space Management 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 112

"maxcandprocs" on page 108

Manually starting threshold migration

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

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 123

Selective file migration

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

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

Migrating selectively 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 process can be quicker if you migrate only premigrated files by using the **stubpremigrated** option.

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

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

An eligible file must meet the following management class requirements:

• The management class assigned to the file enables selective migration.

• A current backup version of the file must exist on your migration server if the management class requires one.

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

"Migrating selectively using the dsmmigrate command" on page 66 **Related reference**:

"dsmmigrate" on page 149

"hsmgroupedmigrate" on page 102

Migration of 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 52

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 Tivoli Storage Manager server or to different 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 Tivoli Storage Manager for Space Management 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 Tivoli Storage Manager server, or if you are using different server stanzas for backup and migration, the Tivoli Storage Manager for Space Management 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, "Tivoli Storage Manager for Space Management client 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.

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

Archiving migrated or premigrated files

You can archive a file that was migrated or premigrated.

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

If you archive a premigrated file, the file is sent from your local file system to Tivoli Storage Manager storage.

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

Retrieving archived files

When you retrieve an archived file, the file status is determined by the **restoremigstate** option in the dsm.opt file.

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

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

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

Back up before migration

You must back up your migrated files to guard against data loss. Use a management class to specify if your files should be backed up before 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 Tivoli Storage Manager for Space Management 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 Tivoli Storage Manager for Space Management 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.

Restoring migrated files

Use the backup-archive client to restore a complete file after the file was migrated.

A file must be backed up and migrated to the same Tivoli Storage Manager server.

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

You can use the Tivoli Storage Manager for Space Management client**dsmmigundelete** command to restore stub files.

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

For more information about the **restore** command, see the *IBM Tivoli Storage Manager for UNIX and Linux Backup-Archive Clients Installation and User's Guide*. If you set the **restoremigstate** option to N0, you restore a complete file to the local file system. The file has resident status. The migrated copy of the file is removed from Tivoli Storage Manager storage when the file expires.

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

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

Related concepts:

"Re-creating stub files with the dsmmigundelete command" on page 72

Related reference:

"restoremigstate" on page 113

"migfileexpiration" on page 110

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

The **dsmmigundelete** command has the following limitations:

- The **dsmmigundelete** command creates a stub file that contains the necessary information to recall the corresponding file from storage. The stub file does not contain any leading bytes of data from the file.
- The recall mode that you previously set for a migrated file is not stored in a re-created stub file. The recall mode for the file is set to normal.
- The Tivoli Storage Manager for Space Management client does not create a stub file if a directory path does not exist in your local file system for a migrated file.
- The Tivoli Storage Manager for Space Management client creates a stub file with the name of the file at the time it was migrated. If you rename a file after it is migrated, the file name is not updated on the server.
- If you have more than one migrated file in storage with the same name that is marked for expiration, the Tivoli Storage Manager for Space Management client creates a stub file again for the file with the most recent modification time (mtime).
- The Tivoli Storage Manager for Space Management client creates a stub file only if the modification time for the migrated file is newer than file on your local file system. If a file is renamed after it is migrated, the Tivoli Storage Manager for Space Management client creates a stub file regardless of the modification time.

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

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

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

Table 17. dsmmigundelete command actions without the expiring option

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

If you issue the **dsmmigundelete** command with the **expiring** option, the Tivoli Storage Manager for Space Management client completes the following action for the file system that you specify:

Table 18. dsmmigundelete command actions with the expiring option

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

Related reference:

"dsmmigundelete" on page 152

Restoring file systems overview

If you restore many stubs with complete files, the file system can run out of space. If the file system runs out of space, the Tivoli Storage Manager for Space Management client must migrate files to make room for more restored files. The restore process can be slowed.

Note:

- When you restore an entire file system, complete 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.

If the backup versions and migrated files reside on the same Tivoli Storage Manager server, you can restore backup versions of resident files but restore migrated and premigrated files as stub files. Use the **dsmc restore** command and set the **restoremigstate** option to YES (the default) in your dsm.sys file.

Related tasks:

"Restoring a file system backed up and migrated to the same server" on page 74 "Restoring a file system backed up and migrated to a different server" on page 74

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 Tivoli Storage Manager for Space Management 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 42

Related reference:

"dsmmigfs add and update" on page 131

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.*

- 1. If needed, reinstall the backup-archive client, API, and Tivoli Storage Manager for Space Management 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.

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. For example, to restore the directory structure for the /home file system, issue the following command:

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

6. Enter the Tivoli Storage Manager for Space Management 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 Tivoli Storage Manager for Space Management 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 N0 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:

"Re-creating stub files with the ${\tt dsmmigundelete}$ command" on page 72

```
Related tasks:
```

"Adding space management to file systems" on page 42

```
Related reference:
```

"dsmmigfs add and update" on page 131

"dsmmigundelete" on page 152

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*.
- Restore the directory structure. For each Tivoli Storage Manager server that manages the file system, issue the following command: dsmc restore -dirsonly -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.
- 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 52

Related reference:

"dsmmigfs addmultiserver, querymultiserver, and removemultiserver" on page 135

"dsmmigundelete" on page 152

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
- Tivoli Storage Manager for Space Management 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 Tivoli Storage Manager for Space Management 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 Tivoli Storage Manager for Space Management client with customized options files. The command line for those clients is sufficient.

After you restore your operating system and are ready to restore you 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 74. If you back up and migrate files to different servers see "Restoring a file system backed up and migrate files to a different server" on page 74.

Related tasks:

"Restoring a file system backed up and migrated to the same server" on page 74 "Restoring a file system backed up and migrated to a different server" on page 74

Chapter 7. Recalling migrated files

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

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

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

Related concepts:

Chapter 1, "Tivoli Storage Manager for Space Management client overview," on page 1

"Recalling migrated files overview" on page 4

Related reference:

"dsmattr" on page 120

Transparent recall processing

Transparent recall processing automatically returns a migrated file to its originating local file system when you access it. Once recalled, the Tivoli Storage Manager for Space Management 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 Tivoli Storage Manager for Space Management client recalls all of your files that are stored on the disk storage

device first. It next recalls all of your files that are stored on the tape storage device. The access time (atime) changes to the current time when you selectively recall a migrated file.

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

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 4

```
Related reference:
```

"dsmrecall" on page 156

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 Tivoli Storage Manager for Space Management client provides a partial file recall mode to recall a portion of a migrated file. This avoids having to recall an entire, potentially large file, when only a small portion of the file is required by an application. This recall mode is only for read access on migrated files that were transferred without compression. Files that were migrated 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 Tivoli Storage Manager for Space Management client calculates which portion of the file to recall based on the offsets contained in the read request. This results in time and disk space savings, because only a portion of the file is recalled.

When you use partial file recall and need to work with a large portion of a large file that has been migrated, it is possible to recall a specified portion of the file in one operation. This minimizes the number of partial file recall requests and access to the server storage based on the **minpartialrecallsize** option value. Use the **dsmrecall** command with the **offset** and **size** options to specify the data range within the file.

For example, to recall 200 MB of the file /usr/cam/video2, starting at 400 MB from the beginning of the file, issue the following command: dsmrecall -offset=400m -size=200m /usr/cam/video2 You can use the following methods to specify which files HSM should recall using partial file recall:

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

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

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

Related concepts:

"Recalling migrated files overview" on page 4

```
Related reference:
```

"dsmmigfs add and update" on page 131

```
"dsmattr" on page 120
```

"dsmrecall" on page 156

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 θ, normal or streaming recall mode takes precedence.

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

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

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

- Change the recall mode
- Modify the file
- Recall the file selectively
- Restore the file
- Set the **minstreamfilesize** option of the **dsmmigfs add** or **dsmmigfs update** command to specify a number to enable or disable an asynchronous recall of migrated files.

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

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

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

Related concepts: "Recalling migrated files overview" on page 4 Related reference: "dsmattr" on page 120 "dsmmigfs add and update" on page 131

How HSM determines which recall mode to use

The Tivoli Storage Manager for Space Management client determines the recall mode depending on the configured recall mode and whether a file is modified.

Table 19 displays the recall mode that is used. The recall mode depends on the following factors:

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

Table 19. Tivoli Storage Manager for Space Management client determining recall modes

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

Setting the recall mode 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 120

Optimized tape recall processing

When files reside on tape, the Tivoli Storage Manager for Space Management client can order a list of files to recall the files efficiently. Optimized tape recall processing is a feature of the selective recall function.

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

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

Optimized tape recall processing provides the following benefits:

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

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

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

Contention for tape resources

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

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

Related concepts:

"Recalling migrated files overview" on page 4

Related reference:

"dsmrecall" on page 156

"hsmmaxrecalltapedrives" on page 107

List files for optimized tape recalls

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

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

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



Figure 2. Ordered list file generation

Input file lists

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

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

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

Generated list files

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

Tape and disk list files

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

- The tape list file has a name with this format: filelist.ordered.tape.tapeID.
- Each entry in the tape file has this format:

order_ID extObjID file_system_ID /filesystem/path/filename

• Files in these lists are ordered by their position on the tape.

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

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

Collection files

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

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

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

Entries in the collection file have this format:

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

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

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





Related concepts:

Chapter 11, "Options files reference," on page 95

Related reference:

"dsmrecall" on page 156

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 156

Chapter 8. Reconciling file systems

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

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

Related concepts:

Chapter 1, "Tivoli Storage Manager for Space Management client overview," on page 1

"Reconciling file systems overview" on page 5

Automatic reconciliation

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

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

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 options 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, "Options files reference," on page 95.

Table 20 describes the tasks that automatic reconciliation completes for files and file systems.

Table 20. Automatic reconciliation tasks

File type	Reconciliation Tasks
Migrated files	• Verifies that a stub file exists on your local file system for each migrated file in storage
	 Marks a migrated file for expiration in the following cases:
	- You deleted the stub file from your local file system
	- You recalled the file and modified it
	Removes a migrated file from storage if it expired
	Updates the status file

Table 20.	Automatic	reconciliation	tasks	(continued)
-----------	-----------	----------------	-------	-------------

File type	Reconciliation Tasks
Premigrated files	Verifies that premigrated files are still valid
	• Marks a premigrated file for expiration in the following cases:
	– You deleted the file from your local file system
	– You accessed the file and modified it
	 Removes a copy of a premigrated file from storage if it expired.
	Updates the status file
Stub files	Records the name of any file for which a stub file exists on your local file system, but a migrated file does not exist in storage. The names are recorded in the orphan.stubs file. For more information about orphaned stub files, see "Resolving orphaned stub files" on page 87.
Status file	Updates the following information in the status file:
	Number of premigrated files
	Number of premigrated blocks
	Number of migrated files
	Number of migrated blocks
	Note: Migration and recall processes update status information dynamically. If any other process changes the state of a file, the status file does not reflect the change until reconciliation is run.

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. Manual reconciliation synchronizes your file system and saves space by removing obsolete copies from storage.

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

Important: A value of θ 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 θ , the file cannot be re-created by the **dsmmigundelete** command.

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

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

dsmreconcile -o /home

Note:

 Since the scout daemon also aids in the reconciliation process, it must be running. For more information about the dsmreconcile command, see "dsmreconcile" on page 160. For more information about the scout daemon, see "The scout daemon" on page 90. • If you reconcile several file systems, increase the value on the **reconcileinterval** option in the dsm.sys file to reduce the impact that the dsmreconcile command might have on system performance.

After you run reconciliation, check the orphan.stubs file in the .SpaceMan directory for each file system that you reconciled. The file lists orphaned stub files. For information about resolving orphaned stub files, see "Resolving orphaned stub files."

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 Tivoli Storage Manager for Space Management client records information about these files in the orphan.stubs file during reconciliation.

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

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

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 N0 in your dsm.opt file, and you then restore a backup version of a migrated file, the file becomes a normal, resident file.

Configuring immediate reconciliation

The Tivoli Storage Manager for Space Management client can expire migration copies on the Tivoli Storage Manager server when migrated or premigrated files are deleted from the space-managed file system.

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

If you set MIGFILEEXPIRATION 0, the migration copy on the Tivoli Storage Manager server is deleted immediately. You cannot recover the file on the space-managed file system with the **dsmmigundelete** command.

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

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

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**_*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 complete a two-way reconciliation by using the **orphancheck** and **fileinfo** options.

Related reference:

"dsmreconcileGPFS.pl" on page 170

"dsmreconcile" on page 160

Chapter 9. Space management daemons

The Tivoli Storage Manager for Space Management client uses the space monitor, recall, scout, and watch daemons to manage file systems automatically.

These daemons are included with the Tivoli Storage Manager for Space Management client. The space monitor daemon (**dsmmonitord**), the scout daemon (**dsmscoutd**), and the recall daemon (**dsmrecalld**) manage space on JFS2 and GPFS file systems. The watch daemon (**dsmwatchd**) runs only on GPFS file systems.

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

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

For information about configuring GPFS integration with the Tivoli Storage Manager for Space Management 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 **dsmmonitord** 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 Tivoli Storage Manager for Space Management 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 154

The recall daemon

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

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

A child recall daemon can recall only one file at a time; however, you can run more than one recall daemon at the same time. To set the minimum and maximum number of child recall daemons that you want to run at one time, use the **minrecalldaemons** and **maxrecalldaemons** options in the dsm.sys file. The minimum number of child recall daemons that you can run at the same time is 1. The default is 3. The maximum number of child aemons 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 165

"dsmrecalld" on page 159
```

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 Tivoli Storage Manager for Space Management 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:
"dsmscoutd" on page 163 "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 Tivoli Storage Manager for Space Management client. Therefore, you do not need to create an RSCT peer domain for the Tivoli Storage Manager for Space Management 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 165

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 155
- "dsmrm" on page 162
- "dmkilld" on page 120
- "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 for the following tasks:

- Balance scheduled services for all client nodes
- Specify that your client node can query the server for scheduled work at specific time intervals or wait for the server to contact your client node 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 started, 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 started, 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 Tivoli Storage Manager for Space Management 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. Options files reference

The Tivoli Storage Manager for Space Management client provides system and space management options that you set either in your dsm.sys file or in your dsm.opt file. The values that you set for space management options determine which server your client node contacts for space management services and affect automatic migration, reconciliation, and recall.

Related tasks:

"Editing the options file dsm.opt" on page 28

"Editing the options file dsm.sys" on page 27

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 163

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 Tivoli Storage Manager for Space Management client from your file system, or when you specify the -o parameter with the **dsmreconcile** command. Queries to the scout daemon for migrated and premigrated files are not done. This is the default.

YES

The **dsmreconcile** command checks for orphans. When orphans stubs are located, their names are recorded in .SpaceMan/orphan.stubs. The **dsmreconcile** command queries the scout daemon for all migrated and premigrated files.

The scout daemon must be running if the **checkfororphans** option is set to YES.

Related reference:

"dsmreconcile" on page 160 "dsmscoutd" on page 163

checkthresholds

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

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

Options File

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

•

Syntax

►►—CHECKThresholds—interval—

Parameters

interval

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

Related reference:

"dsmmonitord" on page 154

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 Tivoli Storage Manager for Space Management 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 Tivoli Storage Manager for Space Management client and the backup-archive client.

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

Authorized user: The value for this option overrides the **DSM_LOG** environment variable.

The log file path cannot be a symbolic link.

Options File

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

Syntax

►►—ERRORLOGName—filespec-

Parameters

filespec

The fully qualified path in which to store error log information. If any part of the path you specify does not exist, the Tivoli Storage Manager for Space Management client creates it.

The log file path cannot be a symbolic link.

errorprog

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

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

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

Options File

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

Syntax

ERRORProg—program-name—

Parameters

program-name

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

Some examples:

errorprog /usr/bin/cat >/tmp/tsm_severe_errors.txt
errorprog /usr/bin/perl /root/dsmseverror.pl

hsmdisableautomigdaemons

Use the **hsmdisableautomigdaemons** option to control the start of the **dsmscoutd** 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 **dsmscoutd** and **dsmmonitord** daemons can be disabled. Stop and restart the **dsmwatchd** daemon to

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

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

▶ → HSMDISAB1eautomigdaemons →	NO	
	L _{YES}	

Parameters

NO Specifies that the Tivoli Storage Manager for Space Management client automatic migration daemons **dsmscoutd** and **dsmmonitord** are started. This is the default.

YES

Specifies that the**dsmscoutd** and **dsmmonitord** daemons are not started.

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

Related tasks:

"Stopping the space management daemons" on page 91

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 N0 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 Tivoli Storage Manager for Space Management client does not migrate a file that is less than 2 minutes old.

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

This option applies only during selective migration.

Options File

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

Syntax



Parameters

NO Files that are less than 2 minutes old cannot be migrated. This is the default.

YES

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

Related reference:

"dsmmigrate" on page 149

hsmeventdestroy

The hsmeventdestroy option specifies whether immediate reconciliation occurs.

When you set HSMEVENTDESTROY yes, the dsmreconciled daemon starts on the affected nodes. The **dsmreconciled** daemon synchronizes the expirations on the Tivoli Storage Manager server with deletions on the space-managed file system.

Syntax



Parameters

NO Immediate reconciliation does not occur. A migrated or premigrated file that is deleted on the space-managed file system is expired on the Tivoli Storage Manager server only during a reconciliation of the file system. This is the default.

YES

Immediate reconciliation occurs. When a migrated or premigrated file is deleted on the space-managed file system, the migration copy on the Tivoli Storage Manager server is expired immediately.

hsmextobjidattr

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

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

The two-way orphan check processing identifies orphans on the file system and on the 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 in the client options file dsm.opt.

Syntax



Parameters

- <u>NO</u> 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.

The **hsmgroupedmigrate** option works in LAN-free configurations and when data is transferred over LAN.

Options File

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

Syntax



Parameters

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

YES

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

hsmlogeventflags

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

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

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

Options File

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

Syntax

►►—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.

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 Tivoli Storage Manager for Space Management client creates it.

The log file path cannot be a symbolic link.

hsmlogretention

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

The HSM log is pruned when the first entry is written to the log after a Tivoli Storage Manager for Space Management command or daemon is started. Some Tivoli Storage Manager for Space Management 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 Do not prune the log. The HSM log grows indefinitely. This is the default.

days

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

D or S

Specifies whether to save the pruned entries. You must also specify the *days* parameter. Enter a space or comma to separate the *days* parameter from the *D* or *S* parameter.

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

Example task: Prune the log every 7 days. Discard the pruned log entries.

Example command: hsmlogr 7 d

S Save the HSM log entries when you prune the log. The pruned entries are copied from the HSM log to the dsmhsm.pru file that is in the same directory as the HSM log.

Example task: Prune the log every 100 days. Save the pruned log entries.

Example command: hsmlogr 100,s

Related reference:

"hsmlogmax" on page 104

Related information:

Chapter 12, "HSM for UNIX and Linux client command reference," on page 115

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 each Tivoli Storage Manager for Space Management client in the GPFS cluster.

If you set hsmmultiserver=N0, all of the command options for the multiple-server environment are disabled. When you set hsmmultiserver=N0, 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 for UNIX and Linux client command reference," on page 115

inclexcl

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

Syntax

▶ — inclexcl—*filespec*—

Parameters

filespec

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

Related concepts:

"Incude-exclude file options" on page 32

Related tasks:

"Creating an include-exclude list" on page 34

maxcandprocs

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

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

Options File

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

Syntax

MAXCANDProcs—number-

Parameters

number

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

maxmigrators

The **maxmigrators** option specifies the maximum number of parallel migration sessions that can run in parallel for each file system. Ensure that you have sufficient resources on the server for parallel migration to occur.

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

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

This option can be set by the 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, make a corresponding increase 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 recall files for the client node. If the number of recall daemons that are running at one time is close to the maximum number, increase the value.

For example, if you use an application that opens many migrated files at once, the application can use all available recall daemons. Another process cannot access a migrated file until a recall daemon is available.

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

Options File

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

Syntax

MAXRECAlldaemons—number—

Parameters

number

Specifies the maximum number of recall daemons that can run in parallel to recall files for the client node. The range is 2 - 99. The default is 20.

maxthresholdproc

The **maxthresholdproc** option specifies the maximum number of threshold migration processes that the Tivoli Storage Manager for Space Management client can start at one time. When a file system runs out of space, the HSM client does not verify the maximum number of threshold migration processes that currently are running. It starts threshold migration as part of the demand migration process *regardless* of the number of threshold migration processes in progress.

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

Options File

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

Syntax

▶▶ — MAXThresholdproc — number-

Parameters

number

Specifies the maximum number of automatic threshold migration processes that the Tivoli Storage Manager for Space Management client can start at one time. The range of values is 1 - 99. The default is 3.

migfileexpiration

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

Options File

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

Syntax

▶ → MIGFileexpiration — days-

Parameters

days

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

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

migrateserver

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

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

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

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

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

Options File

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

Syntax

▶ → MIGRateserver — servername-

Parameters

servername

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

Related reference:

"defaultserver" on page 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 131 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 Tivoli Storage Manager for Space Management client uses the current file system stub size as the minimum size for files that can be migrated. For other (non-GPFS) file systems, the Tivoli Storage Manager for Space Management client uses the file system block or fragment size or stub size as the minimum size for files that can be migrated, whichever is larger.

If you specify a non-zero value, for 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 can run in parallel to recall files for the client node.

Options File

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

Syntax

▶ MINRecalldaemons—number-

Parameters

number

Specifies the minimum number of recall daemons that can run in parallel. The range of values is 1 - 99. The default is 3.

reconcileinterval

The **reconcileinterval** option specifies how often the space monitor daemon reconciles your file systems. Depending on the **checkfororphans** option, the reconciliation either expires or deletes obsolete objects on the server and updates the status file, or checks for orphan stub files and makes metadata updates.

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

Options File

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

Syntax

RECOncileinterval—interval—

Parameters

interval

Specifies the number of hours that must elapse between each successive time your file systems are automatically reconciled on your workstation. If you specify a value of θ , 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 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.

Chapter 12. HSM for UNIX and Linux client command reference

You can use commands to do all HSM tasks. You must use the correct syntax and format when entering commands.

When you issue commands and options, follow these rules:

- Do not precede HSM commands with **dsmc**. Each Tivoli Storage Manager for Space Management 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 Tivoli Storage Manager for Space Management client," on page 25

Standard option formats

You can specify options on some commands. You must use a standard format to specify all options.

Follow these guidelines when using 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

The table gives examples of options that do not define a value.

Table 21. Option format examples: options without values

Command		
dsmmigrate	-recursive -detail	/home/user1/file1
dsmmigrate	-rec -det	/home/user1/file1
dsmmigrate	-r -d	/home/user1/file1
dsmmigrate	/home/user1/file1 -	-r -d

The table gives examples of options when a value is defined.

Table 22. Option format examples: options with values

Commano	d
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 -1=20

In version 6.4 and later, the **optionformat** option is not supported. Short option format is not supported. You must specify all options with the standard option format.

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 **dsmnighelp** 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.

Command	Description
dsmdf	Displays space usage information for a file system. For example, to display space usage information for the /home file system, issue the following command: dsmdf /home

Table 23. HSM commands to display file and file system information

See "dsmdf" on page 125 for more information about this command.

Command	Description
dsmls	Lists files in a directory and displays file conditions. For example, to display information about all files in the /home/user1 directory, issue the following command:
	dsmls /home/user1/*
	See "dsmls" on page 128 for more information about this command.
dsmdu	Displays space usage information for files and directories. For example, to display space usage information for each file in the /home/user/proj1 directory and in all of its subdirectories, issue the following command:
	dsmdu -Allfiles /home/user1/proj1
	See "dsmdu" on page 126 for more information about this command.
dsmmigfs query	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:
	dsmmigfs query /home
	See " dsmmigfs query " on page 139 for more information about this command.
dsmmigundelete	Recreates deleted stub files for migrated files and creates stub files for premigrated files if a corresponding original file does not exist on your local file system. The file then becomes a migrated file. For example, to recreate stub files for migrated files in the /home file system that are not marked for expiration (reconciliation was not run since the files were deleted) issue the command: dsmmigundelete /home
	See " dsmmigundelete " on page 152 for more information about this command.

Table 23. HSM commands to display file and file system information (continued)

Client return codes

The Tivoli Storage Manager for Space Management client command-line interface exits with return codes that accurately reflect the success or failure of the operation.

Scripts, batch files, and other automation facilities can use the return code from the command-line interface. For operations that use the Tivoli Storage Manager scheduler, the return codes are shown in the output of the **QUERY EVENT** administrative command. For cases where the return code is not 0, you can examine the dsmerror.log file. For scheduled events, you can examine the dsmsched.log file.

Return codes have the following meanings:

Table 24. An explanation of client return codes

Code	Explanation
0	All operations completed successfully.

Table 24. An explanation of client return codes (continued)

Code	Explanation
4	 The operation completed successfully, but some files were not processed. There were no other errors or warnings. This return code is common. In most cases, files are not processed for the following reasons: The file satisfies an entry in an exclude list. Excluded files generate log entries only during selective backups. The file was in use by another application and could not be accessed by the client. The file changed during the operation to an extent prohibited by the copy serialization attribute.
8	The operation completed with at least one warning message. Review the dsmerror.log file to determine what warning messages were issued and to assess their effect on the operation.
12	The operation completed with at least one error message (except for error messages for skipped files). For scheduled events, the status is Failed. Review the dsmerror.log file to determine what error messages were issued and to assess their effect on the operation. Generally, this return code means that the error was severe enough to prevent the successful completion of the operation. For example, an error that prevents an entire file system from being processed yields return code 12.

The return code for a client macro is the highest return code issued among the individual commands that comprise the macro. For example, suppose that a macro consists of these commands:

```
selective "/home/devel/*" -subdir=yes
incremental "/home/devel/TestDriver/*" -subdir=yes
archive "/home/plan/proj1/*" -subdir=yes
```

If the first command completed with return code 0; the second command completed with return code 8; and the third command completed with return code 4, the return code for the macro is 8.

For information about the **QUERY EVENT** command, see "**QUERY EVENT** (Query scheduled and completed events)" in the Tivoli Storage Manager information center (http://pic.dhe.ibm.com/infocenter/tsminfo/v6r4/index.jsp?topic=/ com.ibm.itsm.srv.ref.doc/r_cmd_event_query.html).

For more information about the copy serialization attribute, see "Copy serialization attribute" in the Tivoli Storage Manager information center (http://pic.dhe.ibm.com/infocenter/tsminfo/v6r4/index.jsp?topic=/com.ibm.itsm.client.doc/c_mgtc_copyser.html).

HSM command summary

Table 25 provides an alphabetical list of the Tivoli Storage Manager for Space Management client commands, a brief description of each command, and the command page number.

Command and location	Description
dmkilld	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 120.

Table 25.	HSM	command	summary	(continued)
-----------	-----	---------	---------	-------------

Command and location	Description
dsmattr	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 120.
dsmautomig	Starts parallel migration sessions for a file system. See "dsmautomig" on page 123.
dsmdf	Displays space usage information for a file system. See "dsmdf" on page 125.
dsmdu	Displays space usage information for files and directories. See " dsmdu " on page 126.
dsmls	Lists files in a directory and displays file state. See "dsmls" on page 128.
dsmmigfs add, update	Adds space management to a file system, or updates space management attributes for a file system. See "dsmmigfs add and update" on page 131.
dsmmigfs addmultiserver, querymultiserver, removemultiserver	Adds space management to a file system, or updates space management attributes for a file system. See "dsmmigfs addmultiserver, querymultiserver, and removemultiserver" on page 135.
dsmmigfs deactivate, reactivate, remove	Deactivates or reactivates space management for a file system, or removes space management from a file system. See "dsmmigfs deactivate, reactivate, and remove" on page 137.
dsmmigfs sdrreset, enablefailover, disablefailover	Manages recovery from partial system failure (GPFS only). See "dsmmigfs sdrreset, enablefailover, and disablefailover" on page 143.
dsmmigfs query	Displays current space management settings for a file system. See "dsmmigfs query" on page 139.
dsmmigfs globaldeactivate, globalreactivate	Deactivates or reactivates space management for a space-managed client node. See "dsmmigfs globaldeactivate and globalreactivate" on page 138.
dsmmigfs rollback	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 "dsmmigfs rollback" on page 142.
dsmmigfs stop, start, restart	Starts or starts the HSM daemons. See "dsmmigfs stop, start, and restart" on page 144.
dsmmigfs takeover	Transfers the HSM management of a file system to a Tivoli Storage Manager for Space Management client node within the same local GPFS cluster. See "dsmmigfs takeover" on page 145.
dsmmighelp	Displays online help for commands. See "dsmmighelp" on page 146.
dsmmigquery	Displays space management information.
	There are many space management and backup-archive client shared options available using the dsmnigquery -o command. See "HSM and backup-archive client dsmnigquery command shared options" on page 148 for a list of these shared options. See " dsmnigquery " on page 147.
dsmmigrate	Moves selected files from your local file system to Tivoli Storage Manager storage. See "dsmmigrate" on page 149.
dsmmigundelete	Recreates deleted stub files. See "dsmmigundelete" on page 152.
dsmmonitord	Starts the space monitor daemon. See "dsmmonitord" on page 154.
dsmq	Displays information, including recall IDs, for all files that are currently queued for recall. See " dsmq " on page 155.
dsmrecal1	Moves selected files from storage to your local file system. See "dsmrecall" on page 156.

Table 25. HSM command summary (continued)

Command and location	Description
dsmrecalld	Starts the recall daemon. See "dsmrecalld" on page 159.
dsmreconcile	Synchronizes the client and server. See "dsmreconcile" on page 160.
dsmrm	Removes a recall process from the recall queue. See "dsmrm" on page 162.
dsmscoutd	Starts, stops and restarts the scout daemon and shows file system information. See " dsmscoutd " on page 163.
dsmsetpw	Changes the Tivoli Storage Manager password for your client node. See "dsmsetpw" on page 164.
dsmwatchd	Only valid on AIX GPFS and Linux x86_64 GPFS
	Manages failover activities for your Tivoli Storage Manager for Space Management client node. See " dsmwatchd " on page 165.

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.

-options-

Syntax

►►—DMKILLD-

Parameters

options

—Help

Displays syntax and options for the command. Do not specify other options when you specify the **help** option.

dsmattr

The **dsmattr** command sets or changes the recall mode for one or more migrated files. The recall mode determines how the Tivoli Storage Manager for Space Management client recalls a migrated file when you access it. You cannot set a recall mode for a resident or a premigrated file. The recall mode that you set for a migrated file remains associated with that file only 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 Tivoli Storage Manager for Space Management client creates it.

-Help

Displays syntax and options for the command. Do not specify other options when you specify the **help** option.

-RECAllmode=value

Sets a recall mode for one or more migrated files. If you do not use the **RECAllmode** option, the Tivoli Storage Manager for Space Management 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

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 Tivoli Storage Manager for Space Management client sets or displays that recall mode for each migrated file in the directory.

You can use wildcard characters to specify a group of files with similar names. You can issue more than one file specification in a command. If you enter several file specifications, separate each specification with one or more blank spaces.

-FIlelist=file_path

Specifies the path of a file that contains a list of files to be processed by this command.

The entries in the list file must adhere to the following rules:

- Each entry is a fully qualified path to a file or a relative path to a file.
- No entry is a directory object.
- Each entry is on a separate line.

The Tivoli Storage Manager for Space Management client ignores any entry that does not adhere to these rules.

The following rules apply to a file list for the Tivoli Storage Manager for Space Management client. These rules differ from the rules for backup-archive client file lists.

- An entry can contain unprintable characters, but cannot contain a carriage return.
- Wildcard characters are allowed.
- A file 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 Tivoli Storage Manager for Space Management client assumes that the quotation marks are not needed and omits them when the Tivoli Storage Manager for Space Management client processes the entry. If beginning and ending quotation marks are required to identify the file, add double quotation marks before the beginning and after the end of the entry.

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

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 Tivoli Storage Manager for Space Management 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 Tivoli Storage Management client.

Multi-byte characters are interpreted as a set of single bytes all containing valid characters. If you are using multi-byte character sets (MBCS) as your language environment, the Tivoli Storage Manager for Space Management 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 reduce space usage on your file system *before* it reaches the high threshold that you set.

Hidden directories and files are included in automatic migration. 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 dsmmonitord daemon cannot perform threshold migration.

Note: The Tivoli Storage Manager for Space Management client does not migrate contents of symbolic links.

Syntax

	↓ .	
DSHAUTUHIU	filesystemspec storage_pool_name::absolute_filesystem_path	-
▶options_	-	> (

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

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 Tivoli Storage Manager for Space Management client creates it.

-Help

Displays syntax and options for the command. Do not specify other options when you specify the **help** option.

Examples

Task	Command
Start threshold migration for all storage pools on all file systems for which space management is active.	dsmautomig
Start threshold migration for all storage pools in the /home file system.	dsmautomig /home
Start threshold migration for all storage pools in the /home and /test1 file systems.	dsmautomig /home /test1
Start threshold migration for the storage pools named silver and gold for the /fs1 file systems.	dsmautomig /silver::/fs1 gold::/fs1

Start threshold migration for all storage pools in the /fs2 file systems and for dsmautomig gold::/fs1 /fs2 the storage pool named gold in the /fs1 file systems.

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

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 Tivoli Storage Manager for Space Management client creates it.

-Help

Displays syntax and options for the command. Do not specify other options when you specify the **help** option.

-Logname=file_path

Specifies the path and file name of the log file to be used by this command. The kinds of events that are logged to the file are specified with the **hsmlogeventflags** option. Error events are not logged to the file.

filesystemspec

The name of the file system for which you want to display information. The default is all file systems to which you added space management. You can specify more than one file system name, and you can use wildcard characters within a file system name. If you specify more than one file system name, separate each name with one or more blank spaces.

Examples

Task	Command	
Display information for all file systems to which you added space management.	dsmdf	
Display information for the /home file system.	dsmdf /home	
Specify the log file to be used by the dsmdf command	dsmdf -Logname=mylogfile dsmdf -detail -Logname=/tmp/dsmdflog /home dsmdf -d -L=mylogfile	

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 dsmls or dsmdu on remotely mounted GPFS file systems might show incorrect values. See Table 26 for an example of this situation.

Cluster A	Cluster B
Tivoli Storage Manager for Space Management client installed	Tivoli Storage Manager for Space Management 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

Table 26. Space management for locally and remotely mounted file systems

Only file systems of the local cluster can be managed by HSM. The local Cluster A file system is space-managed by Cluster A, but cannot be space-managed by Cluster B, although it is remotely mounted on Cluster B. The Data Management Application Programming Interface (DMAPI) interface does not work for remotely mounted file systems, so incorrect values might be reported by **dsmls** or **dsmdu**.

Syntax



Parameters

options

If you do not specify either of the following options, the Tivoli Storage Manager for Space Management client displays the number of 1 KB blocks that the specified directory and each of its subdirectories use.

-Allfiles

Displays the number of 1 KB blocks that each file in the specified directory and each of its subdirectories use.

-ERRORLOGName=file_path

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file dsm.sys or dsm.opt, or as specified by the **DSM_LOG** environment variable. If any part of the path you specify does not exist, the Tivoli Storage Manager for Space Management client creates it.

—Help

Displays syntax and options for the command. Do not specify other options when you specify the **help** option.

-Summary

Displays only the total of 1 KB blocks that the specified directory and its subdirectories use.

directoryspec

The directory for which you want to display information. The default is the current directory and its subdirectories. Use wildcard characters to specify more than one directory. You can issue more than one directory specification in one command. If you issue several directory specifications, separate each name with one or more blank spaces.

Examples

Task	Command
Display space usage information for the current directory and all of its subdirectories.	dsmdu
Display space usage information for the /migfs3/test directory and all of its subdirectories.	dsmdu /migfs3/test
Display space usage information for each file in the /migfs2/test directory and in all of its subdirectories.	dsmdu -a /migfs2/test
Display the total number of 1 KB blocks that the /migfs2/test directory and all of it subdirectories use.	dsmdu -Summary /migfs2/test

dsm1s

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 notation (s) displays in the file state column.

Restriction:

• Running **dsmls** or **dsmdu** on remotely mounted GPFS file systems might show incorrect values. Consider the following example:

Cluster A	Cluster B
Tivoli Storage Manager for Space Management client installed	Tivoli Storage Manager for Space Management 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

Table 27. Space management for locally and remotely mounted file systems

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 Tivoli Storage Manager for Space Management client creates it.

-Help

Displays syntax and options for the command. Do not specify other options when you specify the **help** option.

-Noheader

Omits column headings from the output for this command.

-Recursive

Displays information about files in subdirectories of the directory.

filespec

The path name for the files that you want to list. The default is all files in the current directory. Use wildcard characters to specify a group of files or all the files in a directory. You can enter more than one file specification in a command. If you enter several file specifications, separate each specification with one or more blank spaces.

-FIlelist=file_path

Specifies the path of a file that contains a list of files to be processed by this command.

The entries in the list file must adhere to the following rules:

- Each entry is a fully qualified path to a file or a relative path to a file.
- No entry is a directory object.
- Each entry is on a separate line.

The Tivoli Storage Manager for Space Management client ignores any entry that does not adhere to these rules.

The following rules apply to a file list for the Tivoli Storage Manager for Space Management client. These rules differ from the rules for backup-archive client file lists.

- An entry can contain unprintable characters, but cannot contain a carriage return.
- Wildcard characters are allowed.
- A file 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 Tivoli Storage Manager for Space Management client assumes that the quotation marks are not needed and omits them when the Tivoli Storage Manager for Space Management client processes the entry. If beginning and ending quotation marks are required to identify the file, add double quotation marks before the beginning and after the end of the entry.

The following 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.	dsmls
List all files in the /migfs2/test directory.	dsmls /migfs2/test
List all files in the /migfs2/test directory and in its subdirectories.	dsmls -Recursive /migfs2/test
List all files whose names begin with tf in the /migfs2/test directory.	dsmls /migfs2/test/tf*
List all files in the list file named /tmp/filelist.	dsmls -filel=/tmp/filelist

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 Tivoli Storage Manager for Space Management 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.

Before you run dsmmigfs add *filesystem* on AIX GPFS and Linux x86_64 GPFS file systems, ensure that the file system is mounted and enabled for Data Management Application Programming Interface (DMAPI) management. Issue the following commands:

For AIX GPFS and Linux x86_64 GPFS: /usr/lpp/mmfs/bin/mmlsfs DevicePath -z

If it is required, change the value as follows: /usr/lpp/mmfs/bin/mmchfs DevicePath -z yes

Run only one **dsmmigfs** command within the local GPFS node group at the same time.

You cannot add space management to your root (/), /tmp, /usr, or /var file systems.

JFS2 Encrypted File System is not supported. Adding space management to a JFS2 EFS with the **dsmmigfs add** command yields an error message.

Note: You can query only information about mounted file systems. If a file system is space-managed, but not mounted, it does not show up within the query command

Syntax



Parameters

Add

Adds space management to your file systems.

Update

Updates one or more space management settings for a file system to which you added space management.

If you change the high and low thresholds or the premigration percentage, the new values take effect immediately. If you change the stub file size, the new size is used for files that are migrated after the change. The size of existing stub files does not change.

If you change the **minmigfilesize** option value, the new value is used the next time a migration candidates list is built. Also, the new value is used only for files that are migrated after the value is changed.

If you change the **readstartsrecall** option value, the new value is used only for files that are migrated after the value is changed.

You can set the quota below the amount of currently migrated and premigrated data. No more files are migrated until automatic and selective recalls reduce the amount of currently migrated and premigrated data to less than the new quota.

filesystemSpec

Specifies a file system name. You can specify more than one file system name, and you can use wildcard characters within a file system name. If you specify more than one file system name, separate each name with one or more blank spaces.

options

Use the option settings that are provided for this command to add or update space management settings.

On GPFS file systems, you can use the GPFS policy engine to monitor space thresholds and search for migration candidates. If you configure the GPFS policy engine to drive automatic migration, the following options for the **dsmmigfs** command have no consequence:

- hthreshold
- lthreshold
- maxcandidates
- maxfiles
- minpartialrecallsize
- minmigfilesize

-HThreshold=n

Specifies the high threshold percentage that you set for space usage on your file systems. Specify a value of 0 -100 percent. The default is 90 percent.

-Lthreshold=n

Specifies the low threshold percentage that you set for space usage on your file systems. Specify a value of 0 through 100 percent. The default is 80 percent.

-ERRORLOGName=file_path

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file dsm.sys or dsm.opt, or as specified by the **DSM_LOG** environment variable. If any part of the path you specify does not exist, the Tivoli Storage Manager for Space Management client creates it.

-LOGname=file_path

Specifies the path and file name of the log file to be used by this

command. The kinds of events that are logged to the file are specified with the **hsmlogeventflags** option. Error events are not logged to the file.

-Maxcandidates=n

Specifies the maximum number of migration candidates the **dsmscoutd** daemon delivers to the automatic migration process during automigration. A value 50 - 1000 is sufficient for good performance. The acceptable range is 9 - 9999999. The default is 10000.

-MAXFiles=n

Specifies the maximum number of files for which the CFI database is sized. The acceptable range is 0 - 4294967295. The default is 0.

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. The value must be greater than the current number of files in the file system. The value must be less than the total likely number of inodes.

You can configure the size of the complete file index (CFI). This value can be queried or set with the maxfiles parameter. If this parameter is set to 0, the CFI allocates the maximum required space. The maximum required space is the maximum number of blocks that are possible in the file system. Less file system space is needed to create the CFI. If the specified value is out of bounds, the nearest boundary is used to allocate the configurable CFI size.

-MINMigfilesize=n

Specifies the minimum size (in bytes) for a file to qualify for migration. The acceptable range is 0 - 2147483647. The default is 0. If you specify 0, the Tivoli Storage Manager for Space Management client uses the current file system stub size plus 1 byte, as the minimum size for files that can be migrated. If you specify a non-zero value, the value must be greater than the current file system stub size. A valid value takes precedence over the global minmigfilesize option setting that is specified in the dsm.sys file.

-MINPartialrecallsize=n (AIX GPFS and Linux x86_64 GPFS only)

Specifies the minimum size (in megabytes) for a file to qualify for partial file recall. The acceptable range is 0 - 999999999. The default value of 0 disables partial file recall for all files.

-MINStreamfilesize=n

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 the Tivoli Storage Manager for Space Management client starts streaming data to the requesting application. Smaller files are recalled in normal recall mode. The recalled portion of the file can be accessed while the file is recalled. The acceptable range is 0 - 999999999. The default is 0.

-Pmpercentage=n

Specifies the percentage of file system space that is available to contain premigrated files. The minimum value is 0. The maximum value is the low threshold. A warning is 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

Specifies 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 999999999999. 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 9999999999999999, the amount of data you can migrate and premigrate is unlimited.

-READEVEnttimeout=n

Specifies the maximum time of inactivity (in seconds) before a streaming recall process times out. The acceptable range is 0 - 999999999. A value of 999999999 means that a streaming recall process does not time out. The default value is 600. This option applies to streaming recall mode only.

-**READStartsrecall**=value

Specifies whether a recall operation starts when an application reads the stub file. Valid values are No and Yes. The default value is No. This option applies only to streaming recall mode.

Regardless of the value of the **readstartsrecall** option, the file is recalled when either of the following conditions are true:

- The recall daemon determines that the stub file does not contain all of the data that is required by a read operation.
- An operation modifies the file.

-SErver=server_name

Overrides the default migration server for this file system. Specify the server to contact for space management services. Define the server in a stanza in your dsm.sys file. If you do not specify a server name, the default migration server that you defined is used. Use a dash (–) to set the server to the default migration server.

-STREAMSeq=n

Specifies the number of megabytes that are buffered before the recall daemon flushes the data to disk. The acceptable range is 0 - 1024. The default value of 0 means that buffering is disabled. This option applies only to streaming recall mode.

-STUbsize=n

Specifies the size of stub files that remain on your local file systems when files are migrated to storage.

For Tivoli Storage Manager for Space Management clients on AIX GPFS and Linux x86_64 GPFS file systems, you can specify 0 or a multiple of the file system block size. The default value is 0.

For Tivoli Storage Manager for Space Management clients on AIX Enhanced Journaled Files Systems (JFS2), you can specify a multiple of the file system fragment size. The default value is the file system fragment size.

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

If HSM is globally deactivated on a node (**dsmmigfs globaldeactivate**), it is reset to active state, if one of the following commands are performed:

dsmmigfs add dsmmigfs remove dsmmigfs update dsmmigfs takeover dsmmigfs rollback dsmmigfs globalreactivate

Examples

Commands are provided for the example tasks.

Task	Command
Add space management to the /hsmmanagedfs1 file system. Set the space to the maximum number of blocks in the file system.	dsmmigfs Add -MAXFiles=0 /hsmmanagedfs1
Update the space management settings for the /hsmmanagedfs2 file system. Set the space (number of blocks) to the specified number.	dsmmigfs update -MAXFiles=10000000 /hsmmanagedfs2
Add space management to the /home file system. Set the high threshold to 80 percent. Set the low threshold to 70 percent. Set the size of stub files to 256 k (KB).	dsmmigfs Add -HT=80 -L=70 -STU=256k /home
Add space management to more than one file system and accept the default values for all space management settings.	dsmmigfs Add /home /test1 /proj*
 Update the space management settings for the /home file system as follows: Change the high threshold to 80 percent. Change the low threshold to 70 percent. Set the size of stub files to 1 megabyte. 	dsmmigfs Update -HT=80 -L=70 -STU=1m /home
Specify the minimum size of files in the /home/user1 file system that are recalled with partial file recall.	dsmmigfs Update -minp=100 /home/user1
Specify that the minimum size for a file that can be migrated from the /home/user2 file system must be at least 1 megabyte.	dsmmigfs Update -minm=1048576 /home/user2
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.	dsmmigfs add -maxfiles=1000000 -lt=50 -ht=80 -pm=10 /myfs
Related tasks:	
"Undating settings from the Tix	voli Storage Manager for Space Management client

"Updating settings from the Tivoli Storage Manager for Space Management client GUI" on page $45\,$

Related reference:

"minmigfilesize" on page 112

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

Adds a Tivoli Storage Manager server to the list of servers that can manage the file system.

QUERYMultiserver

Queries the status of the multiple-server environment. The command prints all 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

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. The current migration and backup processes are not affected. Removal does not affect the recall or restore of files that are migrated or backed up to the Tivoli Storage Manager server.

After you remove 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 you run the dsmRemoveServer.pl script, you can migrate and back up files to another Tivoli Storage Manager server.

filesystemspec

The file system name. Specify only one file system.

options

-ERRORLOGName=file_path

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file dsm.sys or dsm.opt, or as specified by the **DSM_LOG** environment variable. If any part of the path you specify does not exist, the Tivoli Storage Manager for Space Management client creates it.

-Logname=file_path

Specifies the path and file name of the log file to be used by this command. The kinds of events that are logged to the file are specified with the **hsmlogeventflags** option. Error events are not logged to the file.

-SErver=server_name

Specify the target server for the task. This option is valid only in a multiple-server environment. If the file system is managed by multiple servers and you do not specify **server**, the task is attempted with the default migration server.

If a file is coupled with a server, the value of this option must be that coupled server. If you specify another Tivoli Storage Manager server, the task fails.

Related tasks:

"Enabling a file system to be managed by multiple Tivoli Storage Manager servers" on page 52

"Removing a Tivoli Storage Manager server from a multiple-server environment" on page 54

Related reference:

"dsmMultiServerUpgrade.pl" on page 169 "dsmRemoveServer.pl" on page 171

dsmmigfs deactivate, reactivate, and remove

Use the **dsmmigfs** command with the **deactivate**, **reactivate**, or **remove** parameter to deactivate, reactivate, or remove space management from a file system.

You must have root user authority to use this command.

Syntax



Parameters

Deactivate

Deactivates space management for a file system. The Tivoli Storage Manager for Space Management client cannot migrate or recall files, and cannot reconcile the file system. However, you can update space management settings for your file system, and access resident and premigrated files.

REActivate

Reactivates space management for a file system.

REMove

Removes space management from a file system. If you deactivated space management for your file system, reactivate it before you remove space management. If any orphaned stub files are located, the command fails. To remove space management, resolve all orphaned stub files, and issue the **dsmmigfs** command again.

filesystemspec

The file system name. You can specify more than one file system name, and you can use wildcard characters within a file system name. If you specify more than one file system name, separate each name with one or more blank spaces.

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 Tivoli Storage Manager for Space Management client creates it.

-Logname=file_path

Specifies the path and file name of the log file to be used by this command. The kinds of events that are logged to the file are specified with the **hsmlogeventflags** option. Error events are not logged to the file.

Examples

Task	Command
Deactivate space management for the /home file system.	dsmmigfs Deactivate /home
Reactivate space management for the /home file system.	dsmmigfs REActivate /home
Remove space management from the /home file system.	dsmmigfs REMove /home

dsmmigfs globaldeactivate and globalreactivate

Use the **dsmmigfs** command with the **globaldeactivate** or **globalreactivate** parameter to either deactivate or reactivate space management for a space-managed client node.

You must have root user authority to use this command.

Syntax

►► DSMMIGES	GLOBAL Deactivate		M
Dominian o	GLOBALReactivate	options	

Parameters

GLOBALDeactivate

Deactivates space management for all file systems on your client node. The Tivoli Storage Manager for Space Management client cannot migration or recall any files and cannot reconcile any file system. However, you can update space management settings for file systems, add space management to other file systems, or access resident and premigrated files.

GLOBALReactivate

Reactivates space management for your client node. All file systems to which you added space management return to their previous state, including that which you added while space management was globally deactivated.

options

-ERRORLOGName=file_path

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file dsm.sys or dsm.opt, or as specified by the **DSM_LOG** environment variable. If any part of the path you specify does not exist, the Tivoli Storage Manager for Space Management client creates it.

-Logname=file_path

Specifies the path and file name of the log file to be used by this command. The kinds of events that are logged to the file are specified with the **hsmlogeventflags** option. Error events are not logged to the file.

Examples

Task	Command
Globally deactivate space management for your client node.	dsmmigfs GLOBALDeactivate

Task

Command

Globally reactivate space management for your dsmmi client node.

 ${\tt dsmmigfs} \ {\tt GLOBALReactivate}$

dsmmigfs help

Use the **dsmmigfs** command with the **help** parameter to display the **dsmmigfs** command syntax and options.

Syntax

►►—DSMMIGFS—Help

Loptions-

Parameters

He1p

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 Tivoli Storage Manager for Space Management client creates it.

-Logname=file_path

Specifies the path and file name of the log file to be used by this command. The kinds of events that are logged to the file are specified with the **hsmlogeventflags** option. Error events are not logged to the file.

Examples

Task	Command
Display the syntax of the dsmnigfs command.	dsmmigfs h

dsmmigfs query

Use the **dsmmigfs** command with the **query** parameter to display the current space management settings for a file system.

Syntax



Parameters

Query

Display the current space management settings for the named file system.

options

-Detail

Display detailed HSM settings for the file system. Some settings are displayed only when queried with the **detail** option:

- The maximum number of migration candidates that are identified by the scout daemon
- The minimum partial file recall size
- The minimum streaming recall size
- The minimum size (in bytes) that qualifies a file for migration
- The local status for the distributed recall environment
- The status for the failover environment

For AIX GPFS and Linux x86_64 GPFS file systems only:

The **dsmmigfs query** command without the **detail** option displays only locally managed file systems. Use the **detail** option to display information for all space-managed file systems within the GPFS cluster.

When you specify the **detail** option, the HSM client also displays the node name for each node ID.

-Failover

Display a status overview of the failover environment of all HSM managed cluster nodes. 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. The following status can be displayed:

- Active. The node participates in the failover environment within the local GPFS node set.
- Deactivated by User. You disabled failover with 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 Tivoli Storage Manager for Space Management client creates it.

-Logname=file_path

Specifies the path and file name of the log file to be used by this command. The kinds of events that are logged to the file are specified with the **hsmlogeventflags** option. Error events are not logged to the file.

-Node=node_value -parsable

Display the GPFS status of cluster nodes that are managed by the HSM

client. This choice is valid for AIX GPFS and Linux x86_64 GPFS file systems only. The minimum abbreviation of the option is one character (**-n**=*node_value*). The output displays the following information for each node:

- GPFS node name
- GPFS node ID
- GPFS status. The GPFS status can be arbitrating, active, down, or unknown. The status is displayed when you specify the **detail** option.
- · Recall daemon session ID
- Mount disposition
- Ping recall daemon
- Watch daemon session ID
- HSM status. The HSM status can be active or down.

Specify *node_value* by using one of three formats:

Specify node_value as a node name Example: -node=system1

- Specify node_value as a file that contains a list of node names Example: -node=/usr/tivoli/tsm/data/nodes.list
- **Specify** *node_value* **by using a key word that defines a class of nodes** Specify one of the following node classes:
 - **all** All nodes in the GPFS cluster.
 - clientnodes

All nodes that do not participate in file system administration activities.

managernodes

All nodes in the pool of nodes from which file system managers and token managers are selected.

nonquorumnodes

All non-quorum nodes in the GPFS cluster.

nsdnodes

All NSD server nodes in the GPFS cluster.

quorumnodes

All quorum nodes in the GPFS cluster.

Example: -node=quorumnodes

When you specify the **node** option, you can specify the **parsable** option. The **parsable** option is optional. The **parsable** option displays output in a format that can be easily parsed. The minimum abbreviation of the option is four characters (-pars). The parsable format displays the following attributes in order:

Command name Query type GPFS node name GPFS node ID GPFS status HSM status Recall master session ID Mount disposition

Recall daemon readiness

Watch daemon session ID

Each attribute value ends with a colon.

The following output is an example of the format when you specify the parsable option:

dsmmigfs:queryNode:interceptor:2:active:active:51AC92F500000000:YES:YES: 519EF96D00000000:

filesystemspec

The file system name that displays current space management settings. The default is all space managed file systems.

Examples

Task	Command
Display the space management settings for the /migfs2 file system.	dsmmigfs query /migfs2
Display the space management entries that are made in the log file /tmp/migfslog.	dsmmigfs query -L=/tmp/migfslog /migfs2

Related concepts:

Chapter 3, "Configuring the Tivoli Storage Manager for Space Management client," on page 25

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

Parameters

ROLLback

Transfers the HSM management of file systems to the preferred node if the node is different from the current owner node. Enter this command on the preferred node.

options

-ERRORLOGName=file path

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file dsm.sys or dsm.opt, or as specified by the **DSM_LOG** environment variable. If any part of the path you specify does not exist, the Tivoli Storage Manager for Space Management client creates it.

-Logname=file_path

Specifies the path and file name of the log file to be used by this command. The kinds of events that are logged to the file are specified with the **hsmlogeventflags** option. Error events are not logged to the file.

Examples

Task	Command
Transfer to the preferred node.	dsmmigfs rollback

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 Tivoli Storage Manager for Space Management client can take over from an HSM client that is involved in a partial system failure if the following conditions are true:

- The failing Tivoli Storage Manager for Space Management client node has failover enabled.
- There are one or more additional Tivoli Storage Manager for Space Management client nodes within the same GPFS cluster with failover enabled.
- The space-managed file system is mounted on at least one of these nodes.
- A synchronous time exists on the failing nodes and the client nodes.
- The peer node is online

Syntax



Parameters

SDRReset

Do not use this command during normal operation.

The Tivoli Storage Manager for Space Management 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

Activates the node for failover operations within the GPFS cluster.

DISABLEFailover

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 Tivoli Storage Manager for Space Management client creates it.

-Logname=file_path

Specifies the path and file name of the log file to be used by this command. The kinds of events that are logged to the file are specified with the **hsmlogeventflags** option. Error events are not logged to the file.

Examples

Commands are provided for the example tasks.

Task	Command
Deactivate failover operations on the Tivoli Storage Manager for	dsmmigfs disableFailover
Space Management client node.	

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

Starts all HSM daemons on the local client node. The **dsmwatchd** daemon is not affected.

STOP

Stops all HSM daemons. The **dsmrecall** and **dsmmigrate** processes are stopped. The **dsmwatchd** daemon is not affected.

RESTART

Restarts all HSM daemons. The **dsmrecall** and **dsmmigrate** processes are started. The **dsmwatchd** daemon is not affected.

options

-ERRORLOGName=file_path

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file dsm.sys or dsm.opt, or as specified by the **DSM_LOG** environment variable. If any part of the path you specify does not exist, the Tivoli Storage Manager for Space Management client creates it.

-Logname=file_path

Specifies the path and file name of the log file to be used by this command. The kinds of events that are logged to the file are specified with the **hsmlogeventflags** option. Error events are not logged to the file.

Examples

Task	Command
Start all daemons	dsmmigfs START
Stop all daemons	dsmmigfs STOP
Restart all daemons. For example, let them update the configuration set in your dsm.opt and dsm.sys options files.	dsmmigfs RESTART

dsmmigfs takeover

Use the **dsmmigfs** command with the **takeover** parameter to transfer the HSM management of a file system to another the Tivoli Storage Manager for Space Management 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—

Loptions-

Parameters

TAKEover

The **dsmnigfs** command transfers the HSM management of the specified file system to the Tivoli Storage Manager for Space Management client node on which you invoke this command. The transfer must be initiated on a node within the same local GPFS node set.

filespec

The name of the file system you want to takeover.

options

-ERRORLOGName=file_path

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file dsm.sys or dsm.opt, or as specified by the **DSM_LOG** environment variable. If any part of the path you specify does not exist, the Tivoli Storage Manager for Space Management client creates it.

-Logname=file_path

Specifies the path and file name of the log file to be used by this command. The kinds of events that are logged to the file are specified with the **hsmlogeventflags** option. Error events are not logged to the file.

Examples

Task	Command
Transfer the HSM management of the current directory to the Tivoli Storage Manager for Space Management client node within the same local GPFS node set.	dsmmigfs takeover /home/filesystem

dsmmighelp

The **dsmmighelp** command displays online help topics from which you can select general help for commands or message information.

Syntax

DSMMIGHELP _____

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 Tivoli Storage Manager for Space Management client creates it.

Examples

Task	Command
Display online help for HSM commands.	dsmmighelp

dsmmigquery

The **dsmmigquery** command displays information about migrated files, candidates, and management classes.

You must have root user authority to use this command.

The **dsmnigquery** command displays the following information for one or more file systems:

- Migration candidates list
- Ordered recall list for migrated files
- Available management classes
- Current[®] client and server options
- List of all files in the file system

Output from this command is directed to stdout. Use redirection characters and a file name at the end of the command to redirect the output to a file.

Syntax

►►—DSMMIGQUERY

_____filesystemspec_____options___

Parameters

options

-Detail

Use with the **mgmtclass** option to display information about each available management class. If you do not use this option, the HSM client displays the management class name and a brief description only.

-ERRORLOGName=file_path

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file dsm.sys or dsm.opt, or as specified by the **DSM_LOG** environment variable. If any part of the path you specify does not exist, the Tivoli Storage Manager for Space Management client creates it.

-Help

Displays syntax and options for the command. Do not specify other options when you specify the **help** option.

-Mgmtclass

Displays information about each management class that you can assign to your files.

-Options

Displays the current settings for your client and server options. This option is the default.

-SErver=server_name

Specify the target server for the task. This option is valid only in a multiple-server environment. In a multiple-server environment you must specify the **server** option.

This option is valid with the **sortedall** and **sortedmigrated** options.

-SORTEDA11

Lists all files in the file system in this order: resident files, premigrated files, migrated files. Sorts migrated files in the most efficient order for recall.

-SORTEDMigrated

Lists all files that you migrated from the file system to Tivoli Storage Manager storage in the most efficient order for recall.

filesystemspec

The file system for which you want to display information. The default is the current file system. You can specify more than one file system name, and you can use wildcard characters within a file system name. If you specify more than one file system name, separate each name with one or more blank spaces.

Examples

Task	Command
Display the current settings for the client and server options.	dsmmigquery

Display information about management classes that you can assign dsmmigquery -Mgmtclass -Detail to files on your client node.

Related tasks:

"Updating settings from the Tivoli Storage Manager for Space Management client GUI" on page 45

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 Tivoli Storage Manager for Space Management client and the backup-archive client.

The following is a list of the options that are shared between the Tivoli Storage Manager for Space Management 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

- inclexcl
- include
- include.compression
- lanfreecommethod
- 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.

Tip: On large file systems, selective migration can take a long time. Migration can be quicker if you premigrate files before migrating them. The quickest way to create free space requires two steps:

- 1. Prepare for migration by premigrating files.
- When you need to create free space quickly, use the stubmigrated option to migrate only premigrated files.

If the file exists in a file system with a different server stanza than the last processed file, a new session starts for each migrated file. This action can happen as a result of links from one file system to another.

Note:

- 1. you are prompted either to wait for the medium to mount or to skip the file when the following two conditions are met:
 - The **tapeprompt** option is set to YES in the dsm.opt file.
 - The destination for the migrated file is a storage pool that consists of removable media (such as tape)
- The Tivoli Storage Manager for Space Management client does not migrate contents of symbolic links. Symbolic links are not followed during recursive selective migration.

The first file migrates even if the file size exceeds the quota that you specified for the file system. When migration occurs, the **ddf** command displays zero migrated and premigrated bytes for the file system. If the total number of bytes exceeds the quota after the file migrates, the next file is not migrated.

Syntax



Parameters

options

-Detail

Displays the size and file name for each file that you migrate.

-ERRORLOGName=file_path

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file dsm.sys or dsm.opt, or as specified by the **DSM_LOG** environment variable. If any part of the path you specify does not exist, the Tivoli Storage Manager for Space Management client creates it.

-Help

Displays syntax and options for the command. Do not specify other options when you specify the **help** option.

-Logname=file_path

Specifies the path and file name of the log file to be used by this command. The kinds of events that are logged to the file are specified with the **hsmlogeventflags** option. Error events are not logged to the file.

-Premigrate

A copy of the file is sent to Tivoli Storage Manager storage, and the file remains on the local file system. The file state changes to premigrated. You cannot use the **premigrate** option and the **stubpremigrated** option at the same time.

-Recursive

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 Tivoli Storage Manager for Space Management client does not migrate contents of symbolic links. Symbolic links are not followed during recursive selective migration.

-SErver=server_name

Specify the target server for the task. This option is valid only in a multiple-server environment. If the file system is managed by multiple servers and you do not specify **server**, the task is attempted with the default migration server.

If a file is coupled with a Tivoli Storage Manager server, you must specify that server. If a file is not coupled with a Tivoli Storage Manager server, you can specify any Tivoli Storage Manager server that was added to the multiple-server environment.

-STubpremigrated

Only files that are in premigrated state are migrated. The migration process

yields free space on the file system quickly. File copies exist in Tivoli Storage Manager storage. The migration process only replaces the local files with stub files.

filespec

The path and file name of the files that you want to migrate. This parameter is required. You can use wildcard characters to specify a group of files or all files in a directory. You can enter more than one file specification in one command. If you enter more than one file specification, separate each specification with one or more blank spaces.

-FIlelist=file_path

Specifies the path of a file that contains a list of files to be processed by this command.

The entries in the list file must adhere to the following rules:

- Each entry is a fully qualified path to a file or a relative path to a file.
- No entry is a directory object.
- Each entry is on a separate line.

The Tivoli Storage Manager for Space Management client ignores any entry that does not adhere to these rules.

The following rules apply to a file list for the Tivoli Storage Manager for Space Management client. These rules differ from the rules for backup-archive client file lists.

- An entry can contain unprintable characters, but cannot contain a carriage return.
- Wildcard characters are allowed.
- A file 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 Tivoli Storage Manager for Space Management client assumes that the quotation marks are not needed and omits them when the Tivoli Storage Manager for Space Management client processes the entry. If beginning and ending quotation marks are required to identify the file, add double quotation marks before the beginning and after the end of the entry.

The following 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.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.	dsmmigrate -Recursive -Detail /migfs2/test/dir1

Task	Command
Migrate a file named tf04 from the current directory and display the information. Log results in the file /tmp/miglog.	dsmmigrate -Detail -L=/tmp/miglog tf04
Migrate all files in a filelist named /tmp/filelist. Log results in the file /tmp/miglog.	dsmmigrate -logname=/tmp/miglog -filelist=/tmp/filelist
Use a shell application to build a list of files that are owned by user ibm, and then migrate the files.	find /hsmfilesystem -user ibm -print > /tmp/filelist dsmmigrate -filelist=/tmp/filelist
Premigrate the file /migfs2/test/file1, but do not replace the local file with a stub.	dsmmigrate -p /migfs2/test/file1
Migrate all premigrated files in directory /migfs2/test/dir1/ and in all of the subdirectories	dsmmigrate -R -ST /migfs2/test/dir1

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—

Parameters

options

-Recover

Recreates stub files that have been removed from your file system. Use this option if you need better performance during the recover process.

Important: This option recreates all of the existing stub files that overwrites all premigrated or recalled files. Data is lost from the last migration up to when this option is used! After the **dsmmigundelete -recover** command is issued, there is no way you can check whether the data of the local stub file is different from the contents of the file in the 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

Displays detailed information about which stub file is being recreated to the file system.

Do not use this option with the recover option.

-Expiring

Recreates a stub file for a migrated file if a corresponding stub file does not exist on your local file system, whether the migrated file was marked for expiration or not. Or, it creates a stub file for a premigrated file if a corresponding original file does not exist on your local file system, whether the premigrated file was marked for expiration or not.

Issue the **dsmmigundelete** command with the **expiring** option if you ran reconciliation since the files were deleted.

If you do not use the **expiring** option, the HSM client recreates a stub file for a migrated file if a corresponding stub file does not exist on your local file system and the migrated file was not marked for expiration. Or, it creates a stub file for a premigrated file if a corresponding original file does not exist on your local file system, and the premigrated file was not marked for expiration.

Enter the **dsmmigundelete** command *without* the **expiring** option if you did not run reconciliation since the files were deleted.

-ERRORLOGName=file_path

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file dsm.sys or dsm.opt, or as specified by the **DSM_LOG** environment variable. If any part of the path you specify does not exist, the Tivoli Storage Manager for Space Management client creates it.

-Logname=file_path

Specifies the path and file name of the log file to be used by this command. The kinds of events that are logged to the file are specified with the **hsmlogeventflags** option. Error events are not logged to the file.

-SErver=server_name

Specify the target server for the task. This option is valid only in a multiple-server environment. If the file system is managed by multiple servers and you do not specify **server**, the task is attempted with the default migration server.

If a file is coupled with a server, the value of this option must be that coupled server. If you specify another 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 that were accidentally deleted from the /home file system. Recreate stub files regardless of the expiration status of the migration copies (use the expiring option). Reconciliation was run since the files were deleted. Results are logged in the log file /tmp/undeletelog.	dsmmigundelete -expiring -Logname=/tmp/undeletelog /home
Recreate stub files for migrated files in the /home file system. Recreate stub files only for those migration copies that are not marked for expiration (do not use the expiring option). Reconciliation was not run because the files were deleted. Entries are made in the log file /tmp/undeletelog.	dsmmigundelete -L=/tmp/undeletelog /home
Recreates stub files for migrated files that had been removed from the /trullofs file system.	dsmmigundelete -recover /trullofs

dsmmonitord

The **dsmmonitord** command starts the HSM space monitor daemon if it has stopped. If you issue this command and the space monitor daemon is running, action is not taken.

You must have root user authority to use this command.

Syntax

►► DSMMONITORD _____

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 Tivoli Storage Manager for Space Management client creates it.

—Help

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.

The **dsmq** command displays the following information:

- The recall ID
- The hostname of the host that recalls 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 do not display 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, the original name of the file might display as UNKNOWN. If the connection between the Tivoli Storage Manager for Space Management client and the Tivoli Storage Manager server is slow, the original name of the file might display as UNKNOWN. Issue the **dsmq** command again to display 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

DUNG	
	<i>—options</i>

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 Tivoli Storage Manager for Space Management client creates it.

-Help

Displays syntax and options for the command. Do not specify other options when you specify the **help** option.

Examples

Task	Command

Display the status of recall processes.

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 **dsmls** command.

Syntax



Parameters

options

-Detail

Display the size, path, and file name for each file that you recall.

-ERRORLOGName=file path

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file dsm.sys or dsm.opt, or as specified by the **DSM_LOG** environment variable. If any part of the path you specify does not exist, the Tivoli Storage Manager for Space Management client creates it.

-Help

Displays syntax and options for the command. Do not specify other options when you specify the **help** option.

-OFFset=n

Specify the offset from the beginning of the file of the required data range for partial recall (in bytes, kilobytes, megabytes, or gigabytes). Multipliers (k, m, g, K, M, or G) can be used. The range of acceptable values is 0 - 2147483647. There is no default value. 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), -off=5M (megabytes), -off=2G (gigabytes).

-PREView

Generate list files that are optimized for tape recalls but do not recall the files. You must also specify **filelist** and a file system. The **preview** option is not valid when **filelist** specifies a collection file.

-Logname=file_path

Specifies the path and file name of the log file to be used by this command. The kinds of events that are logged to the file are specified with the **hsmlogeventflags** option. Error events are not logged to the file.

-Recursive

Recall files that match the file specification in a directory and its subdirectories. If you do not use this option, files are recalled only for those directories that you specify.

–RESident

When files are recalled, change the state to resident. If you do not use the **resident** option, the recalled files are in the premigrated state.

Remember: During the next reconciliation of the file system, files in resident state are marked for expiration on the Tivoli Storage Manager server.

-SIZE=n

Specify the size of the required data range for partial recall (in bytes, kilobytes, megabytes, or gigabytes). Multipliers (k, m, g, K, M, or G) can be used. The range of acceptable values is 0 - 4294967295. There is no default value. 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** options. The **recursive**, **detail**, and **filelist** options are not valid with the **size** option.

Examples are: -size=10 (bytes), -size=23k (kilobytes), -size=5M (megabytes), -size=2G (gigabytes).

-SErver=server_name

Specify the target server for the task. This option is valid only in a multiple-server environment. If the file system is managed by multiple servers and you do not specify **server**, the task is attempted with the default migration server.

If a file is coupled with a server, the value of **server** must be that coupled server. If you specify another 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 Tivoli Storage Manager for Space Management client recalls all files that match the specification. If a file matches the specification but it is not migrated, an error message displays. If you enter more than one file specification, separate each specification with at least one blank space.

-FILEList=list_file

Process the files that are listed in *list_file*.

The entries in the list file must adhere to the following rules:

- Each entry is a fully qualified path to a file or a relative path to a file.
- No entry is a directory object.
- Each entry is on a separate line.

The Tivoli Storage Manager for Space Management 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 Tivoli Storage Manager for Space Management client generates tape-optimized list files and a collection file. No files are recalled.

If you do not specify the preview option:

The Tivoli Storage Manager for Space Management client recalls files.

If you also specify *file_system_name*:

The Tivoli Storage Manager for Space Management client performs a tape-optimized recall.

If the list file is a collection file that was generated by dsmrecall:

Recall starts. The list files referenced by the collection are processed in the order defined in the collection file. You must specify the collection file generated by **dsmrecall**; you cannot specify a tape list file or disk list file.

If the list file was not generated by dsmrecall:

The Tivoli Storage Manager for Space Management client first reorders the list to optimize recalls from tape. Then the Tivoli Storage Manager for Space Management client recalls the files.

If you do not specify *file_system_name*:

The Tivoli Storage Manager for Space Management client recalls the files in the list file. The Tivoli Storage Manager for Space Management client does not reorder the list to optimize recalls from tape.

file_system_name

Specify a file system. You must also specify the **filelist** option. The Tivoli Storage Manager for Space Management client optimizes the processing of files stored on tape.

Examples

Task	Command
Recall a single file named /migfs1/test/tf04 and display detailed information.	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
Pocall the energified neutring of a file neurod /mfc1/file10 in	demmacall offect=10M cize=E00M /mfc1/file10

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 128 "dsmmigquery" on page 147

dsmrecalld

The **dsmrecalld** command starts a recall daemon if it is not running.

You must have root user authority to use this command.

Note:

- No action is taken if you issue this command while a recall daemon is running.
- When using the backup-archive client to restore space-managed files, dsmrecalld must be running.

Syntax

DSMRECALLD-

Loptions_

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 Tivoli Storage Manager for Space Management client creates it.

-Help

Displays syntax and options for the command. Do not specify other options when you specify the **help** option.

dsmreconcile

The **dsmreconcile** command synchronizes the file systems on your client node with the Tivoli Storage Manager server that you contact for space management services. Use this command at any time to reconcile one or more file systems. Specify how often to automatically reconcile with the **reconcileinterval** option and how long to keep obsolete copies with the **migfileexpiration** option in the dsm.sys file.

You must have root user authority to use this command.

Note: The **dsmreconcile** command must be found with the PATH variable, or the dsmmonitord daemon cannot reconcile the file system.

Syntax

Parameters

options

-Detail

Prints progress messages.

-ERRORLOGName=file_path

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file dsm.sys or dsm.opt, or as specified by the **DSM_LOG** environment variable. If any part of the path you specify does not exist, the Tivoli Storage Manager for Space Management client creates it.

-FILEINFO

Expires and deletes obsolete copies from the Tivoli Storage Manager server. This is the basic reconciliation task and is the default option.

-FILELIST=file_path

Specifies the path of a file that contains a list of files to be processed by this command.

-Help

Displays syntax and options for the command. Do not specify other options when you specify the **help** option.

-ORPHANCHECK

Checks for orphaned files in the local file system.

The HSM client verifies that valid objects for each local stub file exist in the space management pool of the 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 reconciles with 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 Tivoli Storage Manager for Space Management client space monitor daemon, the scout daemon must be running when you specify this option. If reconciliation is controlled by the GPFS policy engine, the scout daemon is not necessary.

When the reconciliation process is running in orphan-check mode, metadata information is updated on the server. Obsolete copies of files are not expired or deleted on the server.

-PREPTWO

This option prepares a file system for a two-way orphan check. Use this option when reconciliation is controlled by the GPFS policy engine.

When reconciliation is controlled by the GPFS policy engine, migrated and premigrated files are associated with DMAPI attribute **ext0bjId**. If you set **hsmextobjidattr**=YES, the DMAPI attribute **ext0bjId** 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 **ext0bjId** attribute when the files were migrated or premigrated.

The **preptwo** option is valid only if **hsmextobjidattr**=YES.

When you specify the **preptwo** option, do not specify any of the following options:

- fileinfo
- filelist
- orphancheck

-SErver=server_name

Specify the target server for the task. This option is valid only in a multiple-server environment. If the file system is managed by multiple servers and you do not specify **server**, the task is attempted with the default migration server.

If a file is coupled with a server, the value of this option must be that coupled server. If you specify another 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 Tivoli Storage Manager for Space Management client reconciles all file systems on your workstation for which space management is active. If you enter more than one file system name, separate each name with at least one blank space.

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 -orphancheck /migfs1
Start reconciliation for file systems /home and /test1.	dsmreconcile /home /test1

Examples

Task	Command
Prepare the two-way orphan check reconciliation for the /migfs1 file system.	dsmreconcile -preptwo /migfs1
Start the policy-based reconciliation on the files that are listed in the /home/dsn/filelist file list.	dsmreconcile -orphancheck -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 Tivoli Storage Manager for Space Management client creates it.

-Help

Displays syntax and options for the command. Do not specify other options when you specify the **help** option.

Examples

Task	Command
Remove recall ID 10 from the queue.	dsmrm 10
Remove recall ID 5 and recall ID 6 from the queue.	dsmrm 5 6
dsmscoutd

The **dsmscoutd** command starts, stops, and restarts the scout daemon. You can also use the command to start a scan, show future scan times, and show past scan information for file systems.

You must have root user authority to use this command.

Syntax



Parameters

All parameters must be entered as indicated. The parameters must be entered in lowercase. The parameters have no abbreviation. The parameters have no short format.

start

Starts the scout daemon. This option is the default, when no parameters are specified.

stop

Stops the scout daemon.

restart

Stops and restarts the daemon.

scanplan

Displays information about the next scan time, include the remaining time until the next scan, for one or more file systems. If the results of this command do not show the status of space-managed file systems, the dsmscoutd daemon lost its communication paths. Communications paths can be lost because of unexpected file system issues such as unmounting a file system. Use the **dsmscoutd restart** command to recover the communication paths for the daemon.

a11

Use this option to include all file systems in the specific action.

filesystemspec

The file system name. You can specify more than one file system name, and you can use wildcard characters within a file system name. If you specify more than one file system name, separate each name with one or more blank spaces.

scannow

Starts scanning of one or more file systems.

statistics

Provides statistics about the current (if there is one) and last scan of the file system.

options

-ERRORLOGName=file_path

Specifies the path and file name of the error log file to be used by this command. Output from only this command is logged to the specified file. Other commands and daemons write output to the error log file specified by the **errorlogname** option in options file dsm.sys or dsm.opt, or as specified by the **DSM_LOG** environment variable. If any part of the path you specify does not exist, the Tivoli Storage Manager for Space Management client creates it.

-Help

Displays syntax and options for the command. Do not specify other options when you specify the **help** option.

dsmsetpw

The **dsmsetpw** command changes the 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 your Tivoli Storage Manager server is at version 6.3.3 or later, and if you use an LDAP directory server to authenticate passwords

Use any of the following characters to create a password:

```
a b c d e f g h i j k l m n o p q r s t u v w x y z
A B C D E F G H I J K L M N O P Q R S T U V W X Y Z
O 1 2 3 4 5 6 7 8 9
~ ! @ # $ % ^ & * _ - + = ` | () { } { } [ ] : ; < > , . ? /
```

Passwords are case-sensitive and are subject to more restrictions that can be imposed by LDAP policies.

If your Tivoli Storage Manager server is at version 6.3.3 or later, and if you do not use an LDAP directory server to authenticate passwords

Use any of the following characters to create a password:

a b c d e f g h i j k l m n o p q r s t u v w x y z A B C D E F G H I J K L M N O P Q R S T U V W X Y Z O 1 2 3 4 5 6 7 8 9 ~ ! @ # \$ % ^ & * _ - + = ` | () { } []:; < > , .?/

Passwords are stored in the Tivoli Storage Manager server database and are not case-sensitive.

If your Tivoli Storage Manager server is earlier than version 6.3.3

Use any of the following characters to create a password:

a b c d e f g h i j k l m n o p q r s t u v w x y z A B C D E F G H I J K L M N O P Q R S T U V W X Y Z O 1 2 3 4 5 6 7 8 9 _ - & + . Passwords are stored in the Tivoli Storage Manager server database and are not case-sensitive.

On the command line, enclose a password in single quotation marks (') if the password contains one or more special characters. The quotation marks ensure that the special characters are correctly interpreted as password characters. Without quotation marks, the special characters can be interpreted as shell escape characters, file redirection characters, or other characters that have significance to the operating system.

Syntax

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 Tivoli Storage Manager for Space Management client creates it.

-Help

Displays syntax and options for the command. Do not specify other options when you specify the **help** option.

oldpassword

The current Tivoli Storage Manager password for your client node.

newpassword

The new Tivoli Storage Manager password to set for your client node.

Examples

Change your current Tivoli Storage Manager password from osecret to dsmsetpw osecret nsecret nsecret.

dsmwatchd

Task

The **dsmwatchd** command manages failover activities for your Tivoli Storage Manager for Space Management client node. If failover is active on your Tivoli Storage Manager for Space Management client node, it checks the status of the **dsmrecalld** command, the **dsmmonitord** command, the **dsmscoutd** command, and the **dsmreconciled** 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.

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 Tivoli Storage Manager for Space Management 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

►►—DSMWATCHD _____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 Tivoli Storage Manager for Space Management client creates it.

-Help

Displays syntax and options for the command. Do not specify other options when you specify the **help** option.

Chapter 13. HSM GPFS clients Perl scripts

Perl scripts leverage GPFS policy and invoke Tivoli Storage Manager for Space Management client commands. You can modify the scripts to suit your environment.

The scripts are valid only for HSM General Parallel File System (GPFS) clients.

The scripts assume that you configured GPFS for integration with the Tivoli Storage Manager for Space Management client.

For information about configuring GPFS integration with the Tivoli Storage Manager for Space Management 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 52

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

be-__dsmNextServername.pl—_file_system_name-

Parameters

file_system_name

Specify the file system. Typically, the GPFS policy engine provides the file system name.

dsmreconcileGPFS.pl

Use the dsmreconcileGPFS.pl script to reconcile a GPFS file system.

The script uses the GPFS policy engine to determine which files are migrated from the specified file system to 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 in the /opt/tivoli/tsm/client/hsm/samples/ directory.

Syntax

→—_dsmreconcileGPFS_______file_system_name ______file_system_name ______

Parameters

--tmpdir=temp_dir

Specify a directory for temporary files. By default, the dsmreconcileGPFS.pl script creates temporary files in the /tmp directory. The dsmreconcileGPFS.pl script deletes the temporary files when processing is complete. Use two dashes before the tmpdir parameter.

-servername=server_name

Specify the 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.Use a single dash before the **servername** parameter.

file_system_name

Specify the file system to reconcile. The following tasks are completed for the file system:

- The Tivoli Storage Manager for Space Management 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 160

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 Tivoli Storage Manager for Space Management client determines how much space is required to recall all files to the file system. If there is not enough space, the Tivoli Storage Manager for Space Management client notifies you. To remove the server, make space available and run the script again.

- 2. Reconciles the specified file system with the Tivoli Storage Manager server that was removed. The following reconciliation tasks are completed:
 - The Tivoli Storage Manager for Space Management 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 Tivoli Storage Manager for Space Management 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 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. Use a single dash before the **server** parameter.

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 54

Related reference:

"dsmreconcileGPFS.pl" on page 170

Chapter 14. Troubleshooting the Tivoli Storage Manager for Space Management client

Use the following information to resolve some typical problems with the Tivoli Storage Manager for Space Management client.

Client trace classes for the Tivoli Storage Manager for Space Management client

The Tivoli Storage Manager for Space Management client provides individual and aggregate trace classes. You can use aggregate trace classes for a shortcut to enable many related trace classes by specifying the aggregate trace class name.

The trace classes in table Table 28 are used for diagnosing problems with the Tivoli Storage Manager for Space Management client. To diagnose problems, use the trace class name with the TRACEFLAG options in the dsm.opt file. For the complete set of trace classes, see *IBM Tivoli Storage Manager Problem Determination Guide*.

Trace class	Description	Usage
PID	Displays the process ID on each trace statement. The PID is part of the SERVICE trace aggregate.	Use this trace class to diagnose problems that might involve multiple processes.
SERVICE	Displays general processing information for the client. The SERVICE aggregate includes all trace classes except MEMORY and *DETAIL classes. The SERVICE trace flag can generate a substantial amount of information. Consider using the tracemax option with the SERVICE trace flag.	Use this trace class when the nature of the problem is unknown. If the SERVICE trace flag is used, it is not necessary to specify other trace flags. For HSM daemon tracing, use this flag if the problem can be quickly re-created.
SM	Displays general processing information of the Tivoli Storage Manager for Space Management client. The SM class is part of the SERVICE or SMALL trace aggregate.	Use this trace class to diagnose problems that are related to HSM (hierarchical storage management).
SMALL	Displays information for HSM.	Use this trace class to enable all trace classes that are related to HSM.
SMXDSM	Displays DMAPI processing information. The SMXDSM class is part of the SERVICE or SMALL trace aggregate.	Use this trace class to diagnose problems that involve DMAPI processing information.

Table 28. Trace classes

Table 28. Trace classes (continued)

Trace class	Description	Usage
TID	Displays the thread ID on each trace statement. The TID class is part of the SERVICE trace aggregate.	Use this trace class to diagnose problems that might involve multiple threads of a single process.

Enabling a Space Management client trace

You can trace the Tivoli Storage Manager for Space Management client by adding trace options to the dsm.opt client options file.

- Determine the trace classes that you want to enable for the Tivoli Storage Manager for Space Management client. See "Client trace classes for the Tivoli Storage Manager for Space Management client" on page 173 for a list of trace classes.
- 2. Add trace options to the dsm.opt client options file.
 - a. Specify the trace classes that you want to enable or disable with the **traceflags** option. You can list one or more trace classes. Separate each trace class with a comma and no space. Use a minus sign (-) in front of a trace class to disable the trace. Put all disabled trace classes at the end of the list. For example, if you want to collect a SERVICE trace without the SESSION or SESSVERB classes, then specify the following option:

traceflags service,-session,-sessverb

 b. Specify the location of the trace output file by using the tracefile option. The tracefile option must specify a complete file path. The following example specifies a complete file path:

tracefile /home/spike/trace.out

c. Optional: Set a maximum size for the trace output by using the **tracemax** option.

If you specify the **tracemax** option, the trace information wraps and writes over the beginning of the data when the output reaches the specified size. The end of the trace information is indicated by "END OF DATA". When data wraps, the end of the output contains "Continued at beginning of file". Consider a maximum trace size if you are trying to capture an event that happens at the end of a long process.

If you specify a **tracemax** value of 1001 or greater and the **tracesegsize** option is not specified, then the trace output is split into several files. Each file is the default segment size, 1000 MB.

Specify the maximum size for the trace output in MB. Valid values are 1 - 4294967295. For example, to set a maximum output size of 4096 MB, use the following option: tracemax 4096

d. Optional: Specify a trace segment size.

Splitting the trace into segments allows easier management of large amounts of trace data. With small files, you do not have to compress or use a separate file splitter utility.

Specify the value in MB. Valid values are 1 - 1000.

Each trace file name contains the value of **tracefile** and appends a number, beginning with 1. For example, if you specify tracefile /home/spike/trace.out and tracesegsize 200, then the trace is segmented into separate files of no more than 200 MB each. The files are named /home/spike/trace.out.1, /home/spike/trace.out.2, and so on.

3. Run the operation or restart the HSM daemons to enable the daemon tracing.

The Tivoli Storage Manager for Space Management client creates trace output during the operation.

You can also configure and start tracing from a command line. The following example specifies trace options in command format:

dsmmigrate -traceflags=service,-session,-sessverb -tracefile=/home/spike/ trace.out -tracemax=4096 -tracesegsize=200

Enabling daemon traces

You can stop and start trace on HSM daemons.

To trace an HSM daemon, you must add trace options to the dsm.opt file and restart the daemon. You must remove the trace options after you restart the daemon, or the Tivoli Storage Manager for Space Management client traces another daemon when the other daemon restarts. Use different commands and steps to restart different HSM daemons.

Daemon traces can create large trace files in a short time. Enable traces only for selected daemons. Re-create the problem immediately. Disable the daemon trace after the problem occurs.

Important: The HSM daemon holds the file handle for the trace file. Do not delete the trace file when the daemon is running. Deleting a trace file when tracing is enabled can yield unwanted results.

Remember: If the option hsmdisableautomigdaemons YES is set in the dsm.opt file, the scout (dsmscoutd) and monitor (dsmmonitord) daemons are not running.

Tracing the watch daemon

You can start traces for only the watch daemon (dsmwatchd).

- 1. Stop the HSM daemons (but not the watch daemon) with the following command: /usr/bin/dsmmigfs stop.
- 2. Add trace options to the dsm.opt client options file as described in "Enabling a Space Management client trace" on page 174.
- **3**. Stop and restart the watch daemon.
 - On Red Hat Enterprise Linux Version 6 (RHEL6), stop and restart the watch daemon by using the following commands:
 - initctl stop HSM
 - initctl start HSM
 - On all other systems besides RHEL6, stop the watch daemon with the following command:

kill -15 dsmwatchd_pid

The watch daemon is restarted by the next system **init** process.

- 4. Remove the trace options from the dsm.opt file. If you do not remove the trace options, other daemons are traced when you restart other daemons.
- 5. Start all HSM daemons other than the watch daemon by using the following command: /usr/bin/dsmmigfs start Because the dsm.opt file contains no tracing information, the other daemons are not traced.

The Tivoli Storage Manager for Space Management client captures data from the watch daemon.

You can stop tracing the watch daemon in two steps:

- 1. Verify that the dsm.opt client options file does not contain trace options.
- 2. Stop and restart the watch daemon as described in 3 on page 175.

Tracing the recall daemon

You can start traces for only the recall daemon (dsmrecalld).

- 1. In a GPFS environment, issue the following command: dsmmigfs disableFailover.
- Add trace options to the dsm.opt client options file as described in "Enabling a Space Management client trace" on page 174.
- 3. Stop and restart the recall daemon.
 - a. Stop the recall daemon by using the following command: dmkilld.
 - b. Start the recall daemon by using the following command: dmrecalld.
- 4. Remove the trace options from the dsm.opt file. If you do not remove the trace options, other daemons are traced when you restart other daemons.
- 5. In a GPFS environment, issue the following command: dsmmigfs enableFailover.

The Tivoli Storage Manager for Space Management client captures data from the recall daemon.

You can stop tracing the recall daemon in two steps:

- 1. Verify that the dsm.opt client options file does not contain trace options.
- 2. Stop and restart the recall daemon as described in 3.

Tracing the scout daemon

You can start traces for only the scout daemon (dsmscoutd).

- 1. In a GPFS environment, issue the following command: dsmmigfs disableFailover.
- Add trace options to the dsm.opt client options file as described in "Enabling a Space Management client trace" on page 174.
- 3. Restart the scout daemon by using the following command: dsmscoutd restart.
- 4. Remove the trace options from the dsm.opt file. If you do not remove the trace options, other daemons are traced when you restart other daemons.
- 5. In a GPFS environment, issue the following command: dsmmigfs enableFailover.

The Tivoli Storage Manager for Space Management client captures data from the scout daemon.

You can stop tracing the scout daemon in two steps:

- 1. Verify that the dsm.opt client options file does not contain trace options.
- 2. Restart the scout daemon by using the **dsmscoutd restart** command.

Tracing the monitor daemon

You can start traces for only the monitor daemon (dsmnonitord).

- 1. In a GPFS environment, issue the following command: dsmmigfs disableFailover.
- 2. Add trace options to the dsm.opt client options file as described in "Enabling a Space Management client trace" on page 174.
- 3. Stop and restart the monitor daemon.
 - a. Stop the monitor daemon by using the following command: kill -15 *dsmmonitord_pid*.
 - b. Start the monitor daemon by using the following command: dsmmonitord.
- 4. Remove the trace options from the dsm.opt file. If you do not remove the trace options, other daemons are traced when you restart other daemons.
- 5. In a GPFS environment, issue the following command: dsmmigfs enableFailover.

The Tivoli Storage Manager for Space Management client captures data from the monitor daemon.

You can stop tracing the monitor daemon in two steps:

- 1. Verify that the dsm.opt client options file does not contain trace options.
- 2. Stop and restart the monitor daemon as described in 3.

Log files to aid troubleshooting

When troubleshooting the Tivoli Storage Manager for Space Management client, you can configure the logs that record HSM activity and error messages.

Related reference:

"Logs for HSM activity and error messages" on page 39

Commands for displaying HSM information

Tivoli Storage Manager for Space Management client commands are provided to display HSM statistics and configuration information.

The following commands yield HSM information for troubleshooting. The commands are organized by the type of information that they yield.

Cluster-specific information

Issue these commands on one node in the cluster.

- To display the configuration and state information for the space management file system, issue the following command:
 - /usr/bin/ls _alR /file_system_name/.SpaceMan
- To display the space management file system and cluster configuration, issue the following command:

/usr/bin/dsmmigfs query -detail

• To display the space management failover configuration, issue the following command:

/usr/bin/dsmmigfs query -failover

• To display the configured Tivoli Storage Manager options, issue the following command:

/usr/bin/dsmmigquery -options

 To display details on the assigned Tivoli Storage Manager server management class configuration, issue the following command: /usr/bin/dsmmigquery -mgmt -detail

Node-specific information

Issue these commands on all nodes in the cluster.

- To display the content of the node-specific space management configuration, issue the following command: /usr/bin/ls _alR /etc/adsm/SpaceMan
- To display the HSM and GPFS related settings, issue the following command:

/usr/bin/cat /etc/adsm/SpaceMan/config/DSMNodeSet

• To display the client option settings, issue the following command:

On AIX systems:

/usr/bin/cat /usr/tivoli/tsm/client/ba/bin64/dsm.opt

On Linux systems:

/usr/bin/cat /opt/tivoli/tsm/client/ba/bin/dsm.opt

• To display the administrator option settings, issue the following command:

On AIX systems:

/usr/bin/cat /usr/tivoli/tsm/client/ba/bin64/dsm.sys

On Linux systems:

/usr/bin/cat /opt/tivoli/tsm/client/ba/bin/dsm.sys

• To display the installed Tivoli Storage Manager for Space Management client executable programs, issue the following command:

On AIX systems:

/usr/bin/ls -alR /usr/tivoli/tsm/client/hsm/bin

On Linux systems:

/usr/bin/ls -alR /opt/tivoli/tsm/client/hsm/bin

Process-specific information

Issue these commands to display process information.

- To display all process IDs, issue the following command: /usr/bin/ps –ef |grep dsm
- To display all process stack information, issue the following command:

On AIX systems:

/usr/bin/procstack process_ID

On Linux systems, use one of the following commands: /usr/bin/pstack process ID

or

/usr/bin/gstack process_ID

Commands for displaying GPFS information

General Parallel File System (GPFS) commands are provided to display General Parallel File System (GPFS) information.

The following commands yield GPFS information for troubleshooting. The commands are organized by the type of information that they yield.

Cluster-specific information

Issue these commands on one node in the cluster.

• To display the configuration information for a GPFS cluster, issue the following command:

/usr/lpp/mmfs/bin/mmlscluster

• To display the available file space on a GPFS file system, issue the following command:

/usr/lpp/mmfs/bin/mmdf device

• To display all the GPFS policy information for a specified file system, issue the following command:

/usr/lpp/mmfs/bin/mmlspolicy device -L

• To display the state of the GPFS daemon on all nodes in the cluster, issue the following command:

usr/lpp/mmfs/bin/mmgetstate -a

• To display the cluster configuration data for a GPFS cluster, issue the following command:

/usr/lpp/mmfs/bin/mmlsconfig

• To create a memory dump of the GPFS DMAPI information, issue the following command:

/usr/lpp/mmfs/bin/mmfsadm dump dmapi

Node-specific information

To display the GPFS node configuration, issue the following command on all nodes in the cluster:

/usr/bin/cat /var/mmfs/gen/mmsdrfs

For more information about GPFS commands, see the *GPFS Problem Determination Guide*.

Commands for displaying operating system information

Operating system commands are provided to display operating system information.

To display the version and release level of the operating system, issue the following command:

On AIX systems: oslevel -r

On Linux systems: uname -a

Other tools for collecting information

A guide to collecting information is available from the Tivoli Storage Manager for Space Management client support team.

For more information about gathering data to troubleshoot the Tivoli Storage Manager for Space Management client, see IBM technote 1268553, *Collecting Data for Tivoli Storage Manager: HSM client problems on UNIX*, at http://www.ibm.com/support/docview.wss?uid=swg21268553.

The document contains a perl script named collect_v8.pl, which enables HSM information to be collected automatically.

HSM common problems and solutions

Tivoli Storage Manager for Space Management client common problems are listed. Typical solutions are suggested.

The following table lists common problems and typical solutions.

Table 29. Common HSM problems and resolutions

Problem	Problem source	Solution
No HSM daemons are running.	The configuration in the dsm.opt file or the dsm.sys file is invalid. The error prevents all HSM daemons from starting.	Run any HSM command. The command output describes the failure. Correct the configuration in the dsm.opt file or the dsm.sys file.
The watch daemon (dsmwatchd) is the only active daemon.	 Any of the following conditions can cause this problem: HSM was stopped on the specified node. Failover is disabled on the specified node. The DMAPI service is not running. 	 Try the following solutions: Start the HSM daemons by issuing the HSM command: dsmmigfs start. The daemons might take up to 30 seconds to start running. Enable the failover on the node by issuing the HSM command: dsmmigfs enablefailover Ensure that GPFS is in the active state on all nodes in the cluster. To verify this state, issue the GPFS command: mmgetstate –a
The mount of DMAPI-enabled file systems fails.	The recall daemon does not run.	Ensure that the recall daemon runs by issuing the command: dsmrecalld. The mount of a DMAPI-enabled file system requires at least one recall daemon in the cluster to be running.
The mount of DMAPI-enabled file systems hangs	 There are two possible causes: On one node in the GPFS cluster, there is an orphaned DMAPI session from a recall daemon that failed. The GPFS configuration parameter, enableLowspaceEvents, is set to yes. To see the current value of this parameter, issue the command: mmlsconfig grep enableLowspaceEvents 	 If there is an orphaned DMAPI session, restart the recall daemon: 1. Stop the recall daemon on all the nodes in the cluster. Issue the command: dmkilld. 2. Start the recall daemon by issuing the command: dsmrecalld. The orphaned DMAPI session is cleaned up while the recall daemon starts. If enableLowspaceEvents=yes, change the value and restart the GPFS daemon on all nodes: 1. /usr/lpp/mmfs/bin/mmchconfig enableLowspaceEvents=no 2. /usr/lpp/mmfs/bin/mmshutdown -a 3. /usr/lpp/mmfs/bin/mmstartup -a

Table 29. Common HSM	problems	and resolutions	(continued)
----------------------	----------	-----------------	-------------

Problem	Problem source	Solution
Several Space Management commands end without processing.	The Tivoli Storage Manager for Space Management client cannot access the node configuration in the /etc/adsm/SpaceMan directory. Typically, this condition is caused by an unmounted /etc file system.	Mount the /etc file system.
Several HSM commands fail with the error message: ANS9418W program-name: Cannot access the GPFS SDR for writing. It might be locked, or the var file system might be full. ANS9414E program-name: Unable to create DSMNodeSet in the SDR. Aborting	 The Tivoli Storage Manager for Space Management client uses the /var/mmfs/gen/mmsdrfs GPFS configuration file to store information for the high-availability features of the client. Under some conditions, the Tivoli Storage Manager for Space Management client cannot update the configuration. Access to this file is denied if the /var file system is full. An HSM process locked a file and this process is no longer working. The file remains locked until the HSM process becomes available to unlock this file. GPFS locks this file for the following conditions: The GPFS daemon stopped on some nodes in the cluster or the network connection is broken. If HSM modifies the configuration on a node, GPFS cannot share the new information with other cluster nodes. Due to many simultaneous updates in the configuration, the GPFS cluster configuration is not synchronized. 	 Try one or more of the following actions: Verify that the /var file system is not full. Verify that GPFS is in an active state on all nodes in the cluster. Unlock the HSM-related locks by issuing the command: dsmmigfs sdrreset Synchronize all the GPFS configuration files by issuing the command: mmchcluster -p LATEST

Appendix A. Control files in the .SpaceMan directory

When you add space management to your file systems, the Tivoli Storage Manager for Space Management client creates control files in a hidden directory named .SpaceMan.

Table 30 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 Tivoli Storage Manager for Space Management client.

The HSM client automatically excludes these files from space management. No action on your part isequired to ensure that the files remain on the local file systems. The Tivoli Storage Manager for Space Management client requires these objects for processing. Do not delete or modify the .SpaceMan directory, its contents, or any of the file ownerships and permissions.

Files	Description
ActiveRecallTab	This file contains the active recall table.
config/dmiFSGlobalState	This file contains global and file system information about the Data Management Application Programming Interface (DMAPI) interface.
config/dmiFSGlobalState.pid	This file is a lock file for the dmiFSGlobalState file.
config/DSM.pid	This file is a lock file for the DSMNodeSet and DSMSDRVersion files.
config/DSMNodeSet	This file stores the HSM configuration.
config/DSMSDRVersion	This file stores the number of changes in the HSM configuration.
dmiFSState	This file stores information about the file system.
dsmmigfstab	If this file exists, it is from an older version of the HSM client. It is needed for transition to current version.
hsmfsconfig.pid	This file is a lock file for the local hsmfsconfig.xml file.
hsmfsconfig.xml	This file contains the space management settings for the file system.
logdir/	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.
metadata/	This directory contains the complete file index (CFI) that was created by the scout daemon. The CFI contains file system information.
orphan.stubs	This file records the orphan files that were identified during reconciliation.
reserved/	This directory contains reserved files for demand migration.
multiserver/BasicRuleSet	This file is a rule set for a file system that is managed by multiple Tivoli Storage Manager servers.

Table 30. Control files stored in the . SpaceMan directory

Files	Description
multiserver/serverlist	This file contains the list of Tivoli Storage Manager servers that manage the file system.
ruleset	This file contains the automatically generated rule set for GPFS.
SDR/	If this directory exists, it is from an older version of the Tivoli Storage Manager for Space Management client.
status	This file records space management statistics for the file system.
<pre>tapeOptimizedRecall/</pre>	This directory contains the recall list files that optimize tape access.

Table 30. Control files stored in the . SpaceMan directory (continued)

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 52

"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 (1 GB).

You can set the size of the CFI by specifying a maximum number of files to be space managed. Use the **maxfiles** option with the **dsmmigfs update** command. For example, if the file system block size is 1 KB and you specify **maxfiles**=1000, the CFI size is 1000×1 KB = 1 MB. In this example the CFI is only 1 MB, regardless of the total size of the file system. The value of the **maxfiles** option must be greater than or equal to the number of files in the file system.

.SpaceMan/reserved

When you add space management to a file system, the Tivoli Storage Manager for Space Management client creates reserved files. The reserved files require 1 GB of space for each file system that is space managed.

The total space required can be estimated as 1 MB for base information plus 1 GB for reserved files for each space-managed file system. If **hsmdisableautomigdaemons**=no (the default value), you must also add space for the CFI.

Related reference:

"dsmmigfs add and update" on page 131

Appendix B. Tivoli support information

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

From the IBM Support Portal at 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.

Communities and other learning resources

In addition to product documentation, many forms of assistance are available to help you get started as you deploy and use the Tivoli Storage Manager family of products. These resources can also help you to solve problems that you might have.

You can use forums, wikis, and other social media tools to ask questions, talk to experts, and learn from others.

User groups

Tivoli Global Storage Virtual User Group

Access this user group at http://www.tivoli-ug.org/storage.

This group makes it possible for individuals from many different industries and types of organizations to share information and work directly with the IBM product experts. Local chapters also exist where members meet in person to share experiences and hear from guest speakers.

ADSM.ORG

Access this mailing list at http://adsm.org.

This independently managed Storage Management discussion forum started when Tivoli Storage Manager was known as ADSTAR Distributed Storage Manager (ADSM). The members of this forum have many years of experience with Tivoli Storage Manager in almost every type of IT environment.

To subscribe to the forum, send 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*.

Tivoli Storage Manager community on Service Management Connect

Access Service Management Connect at http://www.ibm.com/developerworks/ servicemanagement. In the Storage Management community of Service Management Connect, you can connect with IBM in the following ways:

- Become involved with transparent development, an ongoing, open engagement between users and IBM developers of Tivoli products. You can access early designs, sprint demonstrations, product roadmaps, and prerelease code.
- Connect one-on-one with the experts to collaborate and network about Tivoli and the Tivoli Storage Manager community.
- Read blogs to benefit from the expertise and experience of others.

• Use wikis and forums to collaborate with the broader user community.

Tivoli Storage Manager wiki on developerWorks[®]

Access this wiki at https://www.ibm.com/developerworks/servicemanagement/ sm/index.html.

Find the latest best practices, white papers, and links to videos and other resources. When you log on, you can comment on content, or contribute your own content.

Tivoli Support Technical Exchange

Find information about upcoming Tivoli Support Technical Exchange webcasts at http://www.ibm.com/software/sysmgmt/products/support/supp_tech_exch.html. Replays of previous webcasts are also available.

Learn from technical experts who share their knowledge and then answer your questions. The sessions are designed to address specific technical issues and provide in-depth but narrowly focused training.

Other social media sites

LinkedIn

You can join groups on LinkedIn, a social media site for professionals. For example:

- Tivoli Storage Manager Professionals: http://www.linkedin.com/ groups/Tivoli-Storage-Manager-Professionals-54572
- TSM: http://www.linkedin.com/groups?gid=64540

Twitter

Follow @IBMStorage on Twitter to see the latest news about storage and storage software from IBM.

Tivoli education resources

Use these education resources to help you increase your Tivoli Storage Manager skills:

Tivoli Education and Certification website

View available education at http://www.ibm.com/software/tivoli/education.

Use the Search for Training link to find local and online offerings of instructor-led courses for Tivoli Storage Manager.

Education Assistant

Access resources at http://publib.boulder.ibm.com/infocenter/ieduasst/ tivv1r0/index.jsp.

Scroll to view the list of available training videos. Recorded product demonstrations are also available on a YouTube channel.

Searching knowledge bases

If a problem occurs while you are using one of the Tivoli Storage Manager family of products, you can search several knowledge bases.

Begin by searching the Tivoli Storage Manager Information Center at http://pic.dhe.ibm.com/infocenter/tsminfo/v7r1. Within the information center, you can enter words, phrases, or message numbers in the **Search** field to find relevant topics.

Searching the Internet

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

- IBM technotes.
- IBM downloads.
- IBM Redbooks publications.
- IBM Authorized Program Analysis Reports (APARs). Select the product and click **Downloads** to search the APAR list.

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.

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 and download the IBM Support Assistant web page at 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 a software problem might be available from the IBM software support website.

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 the product, or a component for which you want to find a fix.
- 2. Click Downloads, and then click Search for recommended fixes.

If you have not previously customized the site:

Click **Downloads** and search for the product.

Receiving notification of product fixes

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

- From the support page at http://www.ibm.com/support/entry/portal/, click Sign in and 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.
- 2. Click Manage all my subscriptions in the Notifications pane.
- 3. Click the Subscribe tab, and then click Tivoli.
- 4. Select the products for which you want to receive notifications and click **Continue**.
- 5. 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.

- 1. Ensure that you have completed the following prerequisites:
 - a. Set up a subscription and support contract.
 - b. Determine the business impact of the problem.
 - c. Describe the problem and gather background information.
- 2. Follow the instructions in "Submitting the problem to IBM Software Support" on page 191.

Setting up and managing support contracts

You can set up and manage your Tivoli support contracts by enrolling in IBM Passport Advantage[®]. The type of support contract that you need depends on the type of product you have.

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: For critical, system-down, or high-severity issues, 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 level	Description
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). Sign in to access IBM Service Requests and enter your information into the problem submission tool.

By telephone

For critical, system-down, or severity 1 issues, 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**.

Appendix C. 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 IBM Tivoli Storage Manager family of products includes the following accessibility features:

- Keyboard-only operation using standard operating-system conventions
- Interfaces that support assistive technology such as screen readers

The command-line interfaces of all products in the product family are accessible.

Tivoli Storage Manager Operations Center provides the following additional accessibility features when you use it with a Mozilla Firefox browser on a Microsoft Windows system:

- · Screen magnifiers and content zooming
- High contrast mode

The Operations Center and the Tivoli Storage Manager Server can be installed in console mode, which is accessible.

The Tivoli Storage Manager Information Center is enabled for accessibility. For information center accessibility information, see "Accessibility features in the information center" (http://pic.dhe.ibm.com/infocenter/tsminfo/v7r1/topic/ com.ibm.help.ic.doc/iehs36_accessibility.html).

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 (http://www.ibm.com/able) for 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. 19-21, Nihonbashi-Hakozakicho, Chuo-ku Tokyo 103-8510, 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 websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk. IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who 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[®] 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.

Privacy policy considerations

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at http://www.ibm.com/privacy and

IBM's Online Privacy Statement at http://www.ibm.com/privacy/details the section entitled "Cookies, Web Beacons and Other Technologies" and the "IBM Software Products and Software-as-a-Service Privacy Statement" at http://www.ibm.com/software/info/product-privacy.
Glossary

This glossary provides terms and definitions for Tivoli Storage Manager, Tivoli Storage FlashCopy Manager, and associated products.

The following cross-references are used in this glossary:

- *See* refers you from a nonpreferred term to the preferred term or from an abbreviation to the spelled-out form.
- *See also* refers you to a related or contrasting term.

For other terms and definitions, see the IBM Terminology website at www.ibm.com/software/ globalization/terminology.

Α

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 mode, 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.

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.

ACK See acknowledgment.

acknowledgment (ACK)

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. See also server storage, storage pool, storage pool volume.

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. See also inactive file system.

active policy set

The activated policy set that contains the policy rules currently in use by all client nodes assigned to the policy domain. See also policy domain, 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. See also backup version, 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 server. See also 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 central scheduler, client schedule, 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. See also client node session, session.

administrator

A person responsible for administrative tasks such as access authorization and content management. Administrators can also grant levels of authority to users.

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, 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.

application client

A program that is installed on a system to protect an application. The server provides backup services to an application client.

archive

To copy programs, data, or files to another storage media, usually for long-term storage or security. See also 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. See also copy group.

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

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.

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 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 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 demand migration, threshold migration.

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.

В

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. See also 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. See also copy group.

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 storage. More than one backup version can exist in storage, but only one backup version is the active version. See also active version, copy group, inactive version.

bind To associate a file with a management class name. See also archive-retention grace period, management class, rebind.

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 daemon.

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 also administrative command schedule, client schedule.

client A software program or computer that requests services from a server. See also server.

client acceptor

A service that serves the Java applet for the web client to web browsers. On Windows systems, the client acceptor is installed and run as a service. On AIX, UNIX, and Linux systems, the client acceptor is run as a daemon.

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. See also administrative session.

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 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-polling scheduling mode

A method of operation in which the client queries the server for work. See also 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, central scheduler, 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. See also client user-options file, 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. See also client system-options file, options file.

closed registration

A registration process in which only an administrator can register workstations as client nodes with the server. See also 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 to be consistent.

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, strings of characters, or binary data from the data being processed and replaces characters with control characters. Compression reduces the amount of storage space that is required for 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 enterprise configuration, managed server, profile.

conversation

A connection between two programs over a session that allows them to communicate with each other while processing a transaction. See also session.

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, 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 destination, primary storage pool, server storage, storage pool, storage pool volume.

D

daemon

A program that runs unattended to perform continuous or periodic functions, such as network control.

damaged file

A physical file in which read errors have been detected.

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. See also database snapshot, full backup.

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, full backup.

data center

In a virtualized environment, a container that holds hosts, clusters, networks, and data stores.

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 applicationprogramming 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.

data store

In a virtualized environment, the location where virtual machine data is stored.

deduplication

The process of creating representative records from a set of records that have been identified as representing the same entities.

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.

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. See also automatic migration, selective migration, threshold migration.

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. See also copy storage pool.

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

- 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.
- 2. 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.

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 recover 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 also policy domain.

DRM See disaster recovery manager.

DSMAPI

See data storage-management application-programming interface.

dynamic serialization

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. See also shared dynamic serialization, shared static serialization, static serialization.

Ε

- **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, subscription.

enterprise logging

The process of sending events from a 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 An occurrence of significance to a task or system. Events can include completion or

failure of an operation, a user action, or the change in state of a process. See also enterprise logging, receiver.

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, real and virtual storage capacities and channel volume, 2 to the power of 60 or 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, from space management, or from both backup and space management.

exclude-include list

See include-exclude list.

expiration

The process by which files, data sets, or objects are identified for deletion because their expiration date or retention period has passed.

expiring file

A migrated or premigrated file that has been marked for expiration and removal from storage. If a stub file or an original copy of a premigrated file is deleted from a local file system, or if the original copy of a premigrated file is updated, the corresponding migrated or premigrated file is marked for expiration the next time reconciliation is run.

extend

To increase the portion of available space that can be used to store database or recovery log information.

extended attribute (EA)

Names or value pairs that are associated with files or directories. There are three

classes of extended attributes: user attributes, system attributes, and trusted attributes.

external library

A collection of drives that is managed by the media-management system other than the storage management server.

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 migrated file, premigrated file, resident 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, database snapshot, incremental backup.

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.

G

GB See gigabyte.

General Parallel File System (GPFS)

A high-performance shared-disk file system that can provide data access from nodes in a clustered system environment. See also information lifecycle management.

gigabyte (GB)

For processor storage, real and virtual storage, and channel volume, 10 to the

power of nine or 1,073,741,824 bytes. For disk storage capacity and communications volume, 1,000,000,000 bytes.

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.

Globally Unique Identifier (GUID)

An algorithmically determined number that uniquely identifies an entity within a system. See also Universally Unique Identifier.

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.

Η

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. See also hierarchical storage management client, recall, storage hierarchy.

hierarchical storage management client (HSM

- client) A client program that works with the server to provide hierarchical storage management (HSM) for a system. See also hierarchical storage management, management class.
- HSM See hierarchical storage management.

HSM client

See hierarchical storage management client.

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. See also 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. See also active version, backup 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. See also include-exclude file.

incremental backup

The process of backing up files or directories, or copying pages in the database, that are new or changed since the last full or incremental backup. See also selective backup.

individual mailbox restore

See mailbox restore.

information lifecycle management (ILM)

A policy-based file-management system for storage pools and file sets. See also General Parallel File System.

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 Internet Protocol 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. See also migration job.

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.

Κ

KB See kilobyte.

kilobyte (KB)

For processor storage, real and virtual storage, and channel volume, 2 to the power of 10 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.

LAN-free data transfer

See LAN-free data movement.

leader data

Bytes of data, from the beginning of a migrated file, that are stored in the file's corresponding stub file on the local file system. The amount of leader data that is stored in a stub file depends on the stub size that is specified.

library

- 1. A repository for demountable recorded media, such as magnetic disks and magnetic tapes.
- 2. A collection of one or more drives, and possibly robotic devices (depending on the library type), which can be used to access storage volumes.

library client

A server that uses server-to-server communication to access a library that is managed by another storage management server. See also library manager.

library manager

A server that controls device operations when multiple storage management servers share a storage device. See also library client.

local

- Pertaining to a device, file, or system that is accessed directly from a user system, without the use of a communication line. See also remote.
- 2. For hierarchical storage management products, pertaining to the destination of migrated files that are being moved. See also remote.

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 volume

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, physical file, physical occupancy.

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. See also physical occupancy.

logical unit number (LUN)

In the Small Computer System Interface (SCSI) standard, a unique identifier used to differentiate devices, each of which is a logical unit (LU).

logical volume

A portion of a physical volume that contains a file system.

logical volume backup

A back up of a file system or logical volume as a single object.

Logical Volume Snapshot Agent (LVSA)

Software that can act as the snapshot provider for creating a snapshot of a logical volume during an online image backup.

loopback virtual file system (LOFS)

A file system that is created by mounting a directory over another local directory, also known as mount-over-mount. A LOFS can also be generated using an automounter.

LUN See logical unit number.

LVSA See Logical Volume Snapshot Agent.

Μ

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. See also 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

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.

managed server

A 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, enterprise configuration, profile, subscription.

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 bind, copy group, hierarchical storage management client, policy set, rebind.

maximum transmission unit (MTU)

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^{\circledast} 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)

For processor storage, real and virtual storage, and channel volume, 2 to the 20th power or 1,048,576 bytes. For disk storage capacity and communications volume, 1,000,000 bytes.

metadata

Data that describes the characteristics of data; descriptive data.

migrate

To move data to another location, or an application to another computer system.

migrated file

A file that has been copied from a local file system to 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 file state, premigrated file, resident file, stub file.

migration

The process of moving data from one computer system to another, or an application to another computer system.

migration job

A specification of files to migrate, and actions to perform on the original files after migration. See also job file, threshold migration.

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 disks at the same time. The mirroring of data protects it against data loss within the database or 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 also absolute mode, modified 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, 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

A logical drive through which volumes are accessed in a sequential access device class. For removable media device types, such as tape, a mount point is a logical drive associated with a physical drive. For the file device type, a mount point is a logical drive associated with an I/O stream. 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.

Ν

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 file server

See network-attached storage file server.

NAS file server node

See NAS node.

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 server. See also non-native data format.

NDMP

See Network Data Management Protocol.

NetBIOS (Network Basic Input/Output System) 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-attached storage file server (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 See NetBIOS.

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. See also native format.

0

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. See also 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. See also client system-options file, client user-options file.

originating file system

The file system from which a file was migrated. When a file is recalled, it is returned to its originating file system.

orphaned stub file

A file for which no migrated file can be found on the 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.

Ρ

packet In data communication, a sequence of binary digits, including data and control signals, that are 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.

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, logical file, physical occupancy.

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 logical file, logical occupancy, physical file.

plug-in

A separately installable software module that adds function to an existing program, application, or interface.

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. See also active policy set, 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, management class.

premigrated file

A file that has been copied to server 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 server storage. Premigrated files occur on UNIX and Linux file systems to which space management has been added. See also file state, migrated file, resident file.

premigrated files database

A database that contains information about each file that has been premigrated to server storage.

premigration

The process of copying files that are eligible for migration to server 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 copy storage pool, server storage, storage pool, storage pool volume.

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 authority, node privilege class, operator privilege class, policy privilege class, storage privilege class, 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, enterprise configuration, managed server.

profile association

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.

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.

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, management class.

recall To copy a migrated file from server storage back to its originating file system using the hierarchical storage management client. See also selective recall.

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 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 ensuring consistency between the original data repository and the larger system where the data is stored for backup. Examples of larger systems where the data is stored for backup are storage servers or other storage systems. During the reconciliation process, data that is identified as no longer needed is removed.

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

For hierarchical storage management products, pertaining to the origin of migrated files that are being moved. See also local.

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 server 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 server storage and modified. See also file state.

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. See also archive.

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 also administrative command schedule, 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. See also volume.

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. See also 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. See also incremental backup.

selective migration

The process of copying user-selected files from a local file system to server storage and replacing the files with stub files on the local file system. See also demand migration, threshold migration.

selective recall

The process of copying user-selected files from server storage to a local file system. See also recall, transparent recall.

serialization

The process of handling files that are modified during backup or archive processing. See also shared dynamic serialization, shared static serialization, static serialization.

server A software program or a computer that provides services to other software programs or other computers. See also client.

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. See also 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 hierarchical storage management client nodes (space-managed files). See also active-data pool, copy storage pool, primary storage pool, storage pool volume, 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 for the duration of the session. See also administrative session.

session resource usage

The amount of wait time, processor time, and space that is used or retrieved during a client session.

shadow copy

A snapshot of a volume. The snapshot can be taken while applications on the system continue to write data to the volumes.

shadow volume

The data stored from a snapshot of a volume. The snapshot can be taken while applications on the system continue to write data to the volumes.

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. The backup-archive client retries the backup or archive operation a number of times; if the file is being modified during each attempt, the backup-archive client will back up or archive the file on its last try. See also dynamic serialization, serialization, shared static serialization, static serialization.

shared library

A library device that is used by multiple storage manager servers. See also library.

shared static serialization

A copy-group serialization value that specifies that a file must not be modified during a backup or archive operation. The client 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 dynamic serialization, serialization, shared dynamic serialization, 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 hierarchical storage management (HSM) client. The HSM client recalls the file to the client node on demand.

space management

See hierarchical storage management.

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. Stanzas are usually separated by blank lines or colons, and each stanza has a name.

startup window

A time period during which a schedule must be initiated.

static 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 backup-archive client cannot back up or archive the file. See also dynamic serialization, serialization, shared dynamic serialization, 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 tailored to a specific environment, combining servers, systems, storage products, networking products, software, and services.

storage hierarchy

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.

storage pool

A named set of storage volumes that is the destination that is used to store client data. See also active-data pool, copy storage pool, primary storage pool, storage hierarchy.

storage pool volume

A volume that has been assigned to a storage pool. See also active-data pool, copy storage pool, primary storage pool, server storage, volume.

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 server storage. It also contains additional information that can be used to eliminate the need to recall a migrated file. See also migrated file, resident 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 server 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 storage environment, the process of identifying the subscribers to which the profiles are distributed. See also enterprise configuration, managed server.

system privilege class

A privilege class that gives an administrator the authority to issue all server commands. See also privilege class.

Т

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 server storage based on the high and low thresholds that are defined for the file system. See also automatic migration, demand migration, migration job, selective migration.

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.

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. See also macro file, script.

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. See also communication method.

transparent recall

The process that is used to automatically recall a migrated file to a workstation or file server when the file is accessed. See also 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.

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 many classical and historical texts.

Unicode-enabled file space

Unicode file space names provide support for multilingual workstations without regard for the current locale.

Universally Unique Identifier (UUID)

The 128-bit numeric identifier that is used to ensure that two components do not have the same identifier. See also Globally Unique Identifier.

Universal Naming Convention (UNC)

The server name and network name combined. These names together identify the resource on the domain.

UTF-8 Unicode Transformation Format, 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. See also UCS-2.

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 mount point

A directory branch of a file system that is defined as a virtual file system. The virtual file system is backed up to its own file space on the server. The server processes the virtual mount point as a separate file system, but the client operating system does not.

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, server storage, storage pool, 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 (VSS)

A set of Microsoft applicationprogramming interfaces (APIs) that are used 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

An operation that restores data from a local snapshot. The snapshot is the VSS backup that resides on a local shadow volume. The restore operation retrieves the data by using a file-level copy method.

VSS Instant Restore

An operation that restores data from a local snapshot. The snapshot is the VSS backup that resides on a local shadow volume. The restore operation retrieves the data by using a hardware assisted restore method (for example, a FlashCopy operation).

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.

workload partition (WPAR)

A partition within a single operating system instance.

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 (WWN)

A 64-bit, unsigned name identifier that is unique.

WPAR See workload partition.

WWN See worldwide name.

Index

Special characters

.SpaceMan directory 46, 183 .SpaceMan/orphan.stubs 96 add 131 candidatesinterval 95 checkfororphans 96 checkthresholds 96 compressalways 97 compression 97 defaultserver 28,98 dmkilld command 120 dsmattr command 120 dsmautomig command 123 start threshold migration 65 dsmdf command 125 dsmdu command 125, 126 dsmls command 127, 128, 156 dsmmigfs ADDMultiserver 135 dsmmigfs DISABLEFailover 143 dsmmigfs ENABLEFailover 143 dsmmigfs QUERYMultiserver 135 dsmmigfs REMOVEMultiserver 135 dsmmiafs command 145 deactivating space management 57 reactivating space management 57 deactivating space management 137 GPFS 135 GPFS failover 143 help 139 query 139 reactivating space management 137 dsmmigfs command 142, 144 adding HSM 42 adding space management 131 updating space management settings 131 dsmmighelp command 146 dsmmigquery shared options 148 dsmmigquery command 147 dsmmigrate command 149 selective migration 66 start selective migration 66 dsmmigundelete command 152 restoring stub files 72 dsmmonitord command 154 space monitor daemon 89 dsmmonitord daemon 123, 160 dsmq command 155 dsmrecall command 156 dsmrecalld command 159, 165 recall daemon 90 dsmreconcile command 160 dsmrm command 162 dsmscoutd command 163, 165 dsmsetpw command 164 dsmwatchd command 165

errorlogname 98 errorprog 99 exclude.compression 97 filelist parameter 156 hsmdisableautomigdaemons 99 hsmdistributedrecall 100 hsmenableimmediatemigrate 101 hsmenableimmediatemigrate option 62 hsmeventdestroy 101 hsmextobjidattr 102 hsmgroupedmigrate 28, 102 hsmlogeventflags 103 hsmlogmax 104 hsmlogname 105 hsmlogretention 105 hsmlogsampleinterval 106 hsmmaxrecalltapedrives 107 hsmmultiserver 107 hthreshold 131 inclexcl 108 include.compression 97 1threshold 131 maxcandidates 131 maxcandprocs 108 maxfiles 131 maxmigrators 109 maxrecalldaemons 109 maxthresholdproc 110 migfileexpiration 110 migrateserver 28, 111 minmigfilesize 112 minpartialrecallsize setting 131 minrecalldaemons 112 minstreamfilesize 131 passwordaccess option 22 pmpercentage 131 preservelastaccessdate option 62 quota 131 reconcileinterval 113 restoremigstate 28, 113 stubsize 131 update 131

Α

accessibility features 193 activating space management 137 adding HSM dsmmigfs command 131 exported file systems 44 multiple servers 54, 135 nested file systems 43 overview 41 to file system 42 WPARs 44 AIX installing HSM 12, 15 AIX GPFS installing HSM 9, 11 upgrading 13

AIX GPFS (continued) uninstalling HSM 14 AIX JFS installing HSM 14 requirements installing HSM 15 archiving file systems with HSM 69, 70 automatic migration 61 by GPFS policy 3, 61, 63 dsmNextServername.pl 170 file prioritization 63 options 64 threshold migration 123 start manually 65 automignonuse 31

В

backing up before migrating 71 file systems with HSM 69 backup-archive client 1 Bourne and Korn shell 35 Bourne shell variables 35

С

C shell 35 C shell variables 35 candidate selection migration 64 candidates file 183 clusters on GPFS adding space management 37 command dsmmighelp 116 command line return codes for operations 117 command line help displaying 116 commands dmkilld 120 dsmattr 120 dsmautomig 123 start threshold migration 65 dsmdf 125 young files are omitted 101 **dsmdu** 125, 126 young files are omitted 101 dsmls 127, 128, 156 dsmmigfs ADDMultiserver 135 dsmmigfs QUERYMultiserver 135 dsmmigfs REMOVEMultiserver 135 dsmmigfs SDRReset 143 dsmmigfs 139, 144 takeover 145 adding or updating 131 deactivating space management 57, 137

watch daemon 91

commands (continued) dsmmigfs (continued) globally deactivating space management 138 help 139 reactivating space management 57, 137 removing space management 137 dsmmigfs command globally deactivating space management 138 globally reactivating space management 138 dsmmigfs rollback 142 dsmmighelp 146 dsmmigquery 147 dsmmigrate 149 selective migration 66 dsmmigundelete 152 restoring stub files 71 dsmmonitord 154 space monitor daemon 89 dsmg 155 dsmrecall 156 dsmrecalld 159, 165 recall daemon 90 dsmreconcile 160 dsmrm 162 dsmscoutd 163 dsmscoutd command 165 dsmsetpw 164 dsmwatchd 165 watch daemon 91 descriptions 118 displaying GPFS information 179 displaying HSM information 177 displaying operating system information 179 dsmdf 116 dsmdu 116 dsmls 116 dsmmigfs 116 adding HSM 42 dsmmigquery 32 dsmmigundelete 116 using 115 using in executables 117 using in shell scripts 117 configuring HSM overview 41 configuring HSM connection to secondary server 29 configuring Tivoli Storage Manager for Space Management client 25 control files HSM 183 creating stub files for premigrated files 152 customer support contacting 190

D

daemon dsmmonitord 123 dsmscout 48 daemons dsmmonitord 154, 160 dsmrecalld 159 dsmscoutd 163 dsmwatchd 165 dsmwatchd daemon 41 overview 89 recall daemon 90 scout daemon 90 space monitor daemon 89 stopping 91 watch daemon 91 deactivating HSM 57 deactivating space management 137 globally 138 defined 4 deleted stub files, restoring 152 demand migration 3, 61, 63 different backup and migration servers 74 disability 193 disable failover 143 displaying space management information 126 space management settings 139 space usage information 126 displaying GPFS information 179 displaying HSM information 177 displaying online command help 116 displaying operating system information 179 displaying option information 26 DMAPI 41 DSM_CONFIG 34 DSM_CONFIG environment variable 34 DSM_DIR environment variable 34 DSM_LOG environment variable 34 dsm.opt editing 26, 28 dsm.sys editing 26 options 27 dsmdf command 116 dsmdu command 116 dsmerror.log 35 dsmls command 116 dsmmigfs command 116 dsmmigfs SDRReset 143 dsmmighelp command 116 dsmmigquery command 32 dsmmigundelete command 116 dsmMultiServerUpgrade.pl 52, 169 dsmNextServername.pl 170 dsmreconcile reconcileinterval 5 dsmreconcileGPFS.pl 170

Ε

editing dsm.opt 26 dsm.sys 26 enable failover 143 enablelanfree 36 encrypted file system 131 enhancements for V7.1 xv

dsmRemoveServer.pl 171

environment 38 environment variables DSM_CONFIG -34 DSM DIR 34 DSM_LOG 34 error log 39 error messages 39 exclude files (LAN-free) 36 exclude list edit 34 exclude options 32 excluding files from backup 32 from migration 32 from space management 32 executable file return codes from 117 expiration 5 exported file systems adding HSM 44

F

failover configuration 29, 56 file list 128 file migration 61 file systems backing up 69 reconcile 5 reconciling automatic reconciliation 85 GPFS policy 88 immediate reconciliation 88 manual reconciliation 86 orphan stub files 87 overview 85 restoring 69 different backup and migration servers 74 from multiple servers 75 same backup and migration server 74 files archiving 69 backing up 69 migrating 69 restoring 69 retrieving 69 fixes, obtaining 190 free space requirement HSM control files 184 from multiple servers 75

G

globally deactivating space management 138 globally reactivating space management 138 glossary 199 GPFS nodes re-configuring the hostname or GPFS node number 44 GPFS policy reconciliation GPFS policy 88 GPFS policy-driven automatic migration 3, 61, 63 GUI HSM agent verification 23 prerequisites 22 using 24

Η

help 116 high threshold setting 131 high threshold for migration 48 HSM 183 HSM activity log 39 HSM control files 183 space requirement 184 HSM error messages log 39 HSM GPFS clients scripts 169 HSM GUI HSM agent verification 23 prerequisites 22 requirements 7 using 24 HSM log 39 HSM settings high and low thresholds 48 maximum number of files 51 migration quotas 50 minimum file size for migration 47 minimum partial file size 51 minimum streaming file size 51 overview 46 premigration 49 stub file size 50 updating command line 45 HSM GUI 45 hsmagent.opt 38 hsmdisableautomigdaemons 52 hsmfsconfig.xml 41 hsmfsconfig.xml file 46 hsmmultiserver 52

I

IBM Support Assistant 189 immediate reconciliation 88 inclexcl option dsm.sys 32 include files (LAN-free) 36 include list edit 34 include options 32 include-exclude file 25 including files for backup 32 for migration 32 inodes 125 installing HSM 7 AIX 12, 15 AIX GPFS 9, 11 AIX JFS 14

installing HSM (continued) installation directory Linux GPFS 16, 18 prerequisites 7 requirements 7 AIX JFS 15 Linux GPFS 18 uninstalling AIX GPFS 14 Linux GPFS 20 upgrade 7 upgrading AIX GPFS 13 Linux GPFS 19 Internet, searching for problem resolution 189, 190

J

JFS2 EFS 131

Κ

keyboard 193 knowledge bases, searching 189 Korn shell variables 35

L

LAN-free 33, 36 prerequisites 36 LAN-free configuration limitations 55 LAN-free data transfer 36 lanfreecommethod 36 lanfreetcpport 36 limitations HSM 7 LAN-free configuration 55 multiple servers 55 server name encryption 55 server node replication 55 Linux GPFS installing HSM 16, 18 upgrading 19 requirements installing HSM 18 uninstalling HSM 20 list file format 156 list of files displaying 147 log HSM activity 39 HSM error messages 39 log files dsmerror.log 35 logdir directory 183 low threshold setting 131 low threshold for migration 48

Μ

manage recovery 143

management class assign to files 30 default 30 displaying 32 manual reconciliation 86 maximum number of files 51 messages 39 migdestination 31 migrated file setting a recall mode 120 migrating demand 3 overview 3 recalling files 4 selective 3 threshold 3 migrating files automatically 123 selectively 149 migration 47 automatic 61, 63 dsmNextServername.pl 170 options 64 by GPFS policy 61 candidate selection 64 demand 61, 63 eligibility criteria 62 high and low thresholds 48 multiple servers overview 67 overview 61 premigration 63 selective 61, 66 selective migration 149 start manually 66 threshold 61, 63 threshold migration 123 start manually 65 migration candidates list 160 displaying 147 migration quotas 50 migration types 61 migrequiresbkup 31 minimum file size 47 minimum file size for migration 47 multiple server dsmreconcileGPFS.pl 170 multiple servers adding HSM 54 dsmMultiServerUpgrade.pl 169 dsmNextServername.pl 170 dsmRemoveServer.pl 171 enabling HSM 52 limitations 55 node replication 56 overview 52 reconciliation 88 removing a server 54 restoring file systems 75 scripts 169

Ν

nested file systems adding HSM 43 NFS servers 44 no-query restore 71 node replication configuration 29 multiple servers 56 normal recall 78 setting for a migrated file 120

0

obsolete files reconciliation 5 online help displaying for commands 116 optimized tape recall 4 list files 81 overview 80 procedure 84 optional setup 30 options 36 candidatesinterval 95 checkfororphans 96 checkthresholds 96 compressalwavs 97 compression 97 defaultserver 28,98 errorlogname 98 errorprog 99 exclude.compression 97 hsmdisableautomigdaemons 99 hsmdistributedrecall 100 hsmenableimmediatemigrate 62,101 hsmeventdestrov 101 hsmextobjidattr 102 hsmgroupedmigrate 28, 102 hsmlogeventflags 103 hsmlogmax 104 hsmlogname 105 hsmlogretention 105 hsmlogsampleinterval 106 hsmmaxrecalltapedrives 107 hsmmultiserver 107 inclexcl 108 include.compression 97 maxcandprocs 108 maxmigrators 109 maxrecalldaemons 109 maxthresholdproc 110 migfileexpiration 110 migrateserver 28, 111 minmiqfilesize 112 minrecalldaemons 112 passwordaccess 22 preservelastaccessdate 62 reconcileinterval 113 restoremigstate 28, 113 displaying 26 dsm.opt 95 dsm.sys 95 enablelanfree 36 exclude 32 include 32 lanfreecommethod 36 lanfreetcpport 36 rollback 142 setting space management 25 standard format 115 ordered recall list displaying 147

orphan stubs 87 orphan.stubs file 183 overview 7, 46, 73

Ρ

packages AIX JFS2 15 partial file recall 4, 78 minimum partial file size 51 partial recall mode setting for a migrated file 120 Passport Advantage 190 password 164 passwordaccess 22 encryption 22 generate 22 managing 22 perl scripts 169 dsmMultiServerUpgrade.pl 169 dsmNextServername.pl 170 dsmreconcileGPFS.pl 170 dsmRemoveServer.pl 171 policies displaying 30, 32 premigration 63 premigration percentage 49, 63 premigration setting 131 premigrdb.dir file 183 premigrdb.pag file 183 prerequisites 36 installing HSM LAN-free 36 problem determination describing problem for IBM Software Support 191 determining business impact for IBM Software Support 191 submitting a problem to IBM Software 191 publications download xi Tivoli Storage Manager xi

Q

quota setting 131 quotas 50

R

reactivating HSM 57 reactivating space management 137 globally 138 recall daemon 90 **dmkilld** command 120 starting 159 recall ID 155 recall mode normal 120 setting for a migrated file 120 recall process displaying status of 155 recalling mode selection 80 recalling (continued) normal recall 78 overview 77 partial file recall 78 selective recall 77 setting the recall mode 80 streaming recall 79 tape-optimized recall list files 81 overview 80 procedure 84 transparent recall 77 recalling files canceling 162 migrated files 4 removing from queue 162 selectively 156 reconcile file systems 5 reconciliation dsmreconcileGPFS.pl 170 policy-based 160 reconciling automatic reconciliation 85 GPFS policy 88 immediate reconciliation 88 manual reconciliation 86 manually 160 orphan stub files 87 overview 85 reconciling file systems 160 reconfiguring HSM GPFS clusters 44 recreating deleted stub files 152 registration HSM client 21 closed registration 21 open registration 21 removing space management 137 removing a server 54 removing HSM 58 GPFS clusters 44 requirements installing HSM 7 AIX GPFS 11 AIX JFS 15 Linux GPFS 18 reset in the SDR 143 resolving 87 restarting space management daemons 144 restoring 73, 74, 75 dsmmigundelete command 72 a disk 76 deleted stub files 152 entire system 76 file systems with HSM 69 no-query restore 71 standard-query restore 71 stub files 72 restrictions HSM 7 retrieving file systems with HSM 70 retrieving files 69 return codes for operations 117 rollback option 142

root user tasks 154 creating stub files for premigrated files 152 recreating deleted stub files 152 starting a scout daemon 163

S

same backup and migration server 74 SAN LAN-free 36 scheduling displaying schedules 94 options 93 overview 93 starting 93 scout daemon 90 starting 163 scripts for HSM GPFS clients 169 perl 169 selected files migrating 149 recalling 156 selecting HSM options 25 selective migration 3, 61, 66 performing 149 start migration 66 selective recall 4, 77, 156 server name encryption limitations 55 server node replication limitations 55 services scheduling displaying schedules 94 options 93 overview 93 starting 93 setting environment variables 34 setting space management options 25 setting the recall mode 80 setting up 25 setup tasks 30 shell scripts 1 return codes from 117 using commands in 117 software support describing problem for IBM Software Support 191 determining business impact for IBM Software Support 191 submitting a problem 191 Software Support contacting 190 space management 25 adding to a file system 131 agent 38 clusters on GPFS 37 files candidates file 183 logdir directory 183 orphan.stubs file 183 premigrdb.dir file 183 premigrdb.pag file 183 status file 183 installing 7

space management (continued) maintaining 131 removing 137 updating a file system 131 space management settings 131 maxcandidates 131 displaying 139 high threshold 131 premigration setting 131 quota 131 stub file size 131 space monitor daemon 89 starting 154 space requirement HSM control files 184 space usage information displaying 126 space usage, displaying 116, 125 hsmenableimmediatemigrate 101 hsmgroupedmigrate 102 young files are omitted 101 spacemgtechnique 31 standard option format 115 standard-query restore 71 starting space management daemons 144 status file 183 stopping space management daemons 144 storage pool 1 dsmautomig command 123 streaming recall 79 minimum streaming file size 51 overview 4 stub file size setting 131 stub files recreating 152 recreating deleted 116 size 50 summary of changes V7.1 xv support contract 190 support information 187 support subscription 190 synchronizing client and server files 85 synchronizing client-server files manual reconciliation 160

Т

tape recall, optimized list files 81 overview 80 procedure 84 tasks 30 threshold migration 3, 61, 63 start manually 65 starting manually 123 thresholds for migration 48 Tivoli Storage Manager for Space Management client overview 1 Tivoli Storage Manager password 164 trace enable client trace 174 enable daemon trace 175 monitor daemon 177 recall daemon 176

trace (continued) scout daemon 176 watch daemon 175 trace classes Tivoli Storage Manager for Space Management client 173 transfer HSM management of file system 142, 145 transparent recall 4, 77 troubleshooting commands displaying GPFS information 179 displaying HSM information 177 displaying operating system information 179 common problems and solutions 180 other tools 180 trace classes 173 truncate migrated files 4 two-way orphan check 160

U

update files in SDR 143 updating space management settings 131 upgrade multiple servers dsmMultiServerUpgrade.pl 169 using commands 115

V

variables 35 variables, environment 34 Version 7.1 changes xv

W

watch daemon 41, 91 WPARs, adding HSM 44

I B M R

Product Number: 5608-E12

Printed in USA