

IBM Tivoli Storage Manager
for Space Management for UNIX



User's Guide

Version 5 Release 2

IBM Tivoli Storage Manager
for Space Management for UNIX



User's Guide

Version 5 Release 2

Note

Before using this information and the product it supports, read the general information in Appendix C, "Notices," on page 163.

Third Edition (December 2003)

This edition applies to version 5, release 2, modification 2 of the IBM Tivoli Storage Manager (5698-ISM), IBM Tivoli Storage Manager Extended Edition (5698-ISX), IBM Tivoli Storage Manager for Storage Area Networks (5698-SAN), and to all subsequent releases and modifications until otherwise indicated in new editions or technical newsletters.

Order publications through your IBM representative or the IBM branch office that serves your locality.

Your feedback is important in helping to provide the most accurate and high-quality information. If you have comments about this manual or any other IBM Tivoli Storage Manager documentation, see "Downloading or ordering publications" on page xi.

© Copyright International Business Machines Corporation 1995, 2003. All rights reserved.

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

Contents

Tables	v
-------------------------	----------

About This Book	vii
----------------------------------	------------

Who Should Read This Manual	viii
IBM Tivoli Storage Manager Web site	viii
Conventions Used in This Manual	viii
Reading Syntax Diagrams	viii
Related Information	xi
Downloading or ordering publications	xi
Contacting customer support	xi
Reporting a problem	xii
Internet	xiii

Chapter 1. Summary of the HSM Client	1
---	----------

Migrating Files from Your Local File System	1
Premigrating Files from Your Local File System	2
Recalling Your Migrated Files	2
Reconciling Your File Systems	3

Chapter 2. Installing the HSM Client.	5
--	----------

Moving Your Files to Another Directory	6
Installing the HSM Client on AIX GPFS File Systems	6
Environment Requirements	8
Installation	9
Installing the HSM Client on AIX JFS File Systems	11
Environment Requirements	11
Installation	12
Obtain VFS and Group Numbers	13
Installing the HSM Client on HP-UX VxFS File Systems	14
Environment Requirements	14
Installation	14
Installing the HSM Client on Linux86 GPFS File Systems	16
Environment Requirements	16
Installation	17
Installing the HSM Client on Solaris VxFS File Systems	19
Environment Requirements	19
Installation	20
Registering Your Workstation with a Server	21
Use Closed Registration	21
Use Open Registration	21
Changing your Password	22
Starting the HSM Client	23

Chapter 3. Selecting Your Options	25
--	-----------

Setting Options for Space Management	25
Display Option Information	25
Edit Your Options Files	26
Performing Optional Setup Tasks	29
Assign Management Classes to your Files	29
Display Information About Management Classes	32
Modify your Include-Exclude Options File	32

Set Environment Variables	34
Set Up LAN-free Data Transfer for HSM on AIX GPFS	35
LAN-free prerequisites	35
LAN-free options	36
Set Up Space Management for AIX in an HACMP Environment	36
Limitations of HSM Support for the HACMP Environment	37
Start and Stop HSM in an HACMP Environment	37
Set Up Space Management for Linux86 in an IBM Cluster 1350 Environment	39
Limitations of HSM Support for the IBM Cluster 1350 Environment	40
Start and Stop HSM in an IBM Cluster 1350 Environment	40

Chapter 4. Managing Your File Systems	43
--	-----------

Mounting Your File Systems First	43
Adding Space Management to Your File Systems	44
Select Space Management Settings	45
Add Space Management	48
Add Space Management to Nested File Systems	49
Managing Your File Systems	50
Update Space Management Settings	50
Deactivate Space Management from Your File Systems	51
Globally Deactivate Space Management	51
Reactivate Space Management on Your File Systems	52
Globally Reactivate Space Management	52
Remove Space Management	53

Chapter 5. Migrating Your Files	55
--	-----------

Migrating Your Files Automatically	56
Premigrate Your Files	56
Start Threshold Migration Manually	57
Build a New Migration Candidates List	57
Migrating Selected Files	58
Migrate Selected Files	59

Chapter 6. Backing Up and Restoring Files	61
--	-----------

Backing Up, Migrating, and Restoring Your Files	61
Back Up Your Files Before You Migrate Them	61
Back Up Your Files After You Migrate Them	62
Back Up and Restore Your Premigrated Files	62
Restore Your Migrated Files	62
Recreating Stub Files for Your Migrated Files	63
Restoring Your File Systems	65
Back Up and Migrate Files to the Same Server	65
Back Up and Migrate Files to a Different Server	66
Restoring A Disk	67
Restoring Your Operating System	67
Restore Your File Systems	67

Backup Server and Migration Server are Different 68

Chapter 7. Archiving and Retrieving Files 69

Archiving Migrated or Premigrated Files 69
Retrieving Archived Copies of Files 70

Chapter 8. Recalling Your Migrated Files 71

Recalling Your Migrated Files Automatically 71
 Recall Your Files in an AIX JFS Environment 71
 Recall Your Files in an AIX GPFS Environment 74
 How HSM Determines Which Recall Mode to Use 75
Recalling Your Migrated Files Selectively 76
 Use Selective Recall. 76
Using the dsmrecall Command 77
 Use the dsmattr Command 77
 Use the dsmmode Command 77

Chapter 9. Reconciling Your File Systems 79

Starting Automatic Reconciliation 79
Reconciling Your File Systems Manually. 80
 Perform All Reconciliation Tasks 80
 Synchronize Your Client and Server 81
Resolving Orphaned Stub Files 81
Reconciling Your File Systems Using Commands 82
Using Space Management Daemons 82
 Start the Space Monitor Daemon 83
 Start the Recall Daemon 83
 Start the Scout Daemon 83
 Stop the Space Management Daemons 84

Chapter 10. Using Utilities 85

Displaying Information About Your File Systems 85
 Display Information from the Command Line 85

Chapter 11. Scheduling Services 87

Setting Scheduling Options 87
Starting the Client Scheduler 87
Displaying Information About Scheduled Services 87

Chapter 12. Using Options 89

Chapter 13. Using Commands 113

Entering Commands 113
 Enter Options With Commands 113
 Use the Standard Format 113
 Use the Short Format. 114
Displaying Online Help for Commands 114
Hierarchical Storage Management Commands 114

Appendix A. Using the Installfsm Program 159

Appendix B. The .SpaceMan Directory Space Management Objects. 161

Appendix C. Notices 163
Trademarks 165

Glossary 167

Index 179

Tables

1.	Description of Syntax Diagram Items	ix	11.	Environment Variables	34
2.	Tivoli Storage Manager Client Manuals	xi	12.	Setting the Age and Size Factors	47
3.	HSM Client Installation Procedures	5	13.	Recall Modes	75
4.	Package Names and Descriptions	18	14.	Reconciliation Tasks.	79
5.	Package Names and Descriptions	20	15.	Examples Using the dsmreconcile Command	82
6.	Setting Space Management Options in Your dsm.sys File	28	16.	Displaying Space Management Information	86
7.	Setting Space Management Options in Your dsm.opt File	29	17.	Space Management Options in Your Options Files	90
8.	Space Management Attributes in a Management Class	31	18.	HSM Commands	114
9.	Assigning Management Classes to Your Files	32	19.	Space Management Option Settings	125
10.	Include and Exclude Statements.	33	20.	HSM Control Files	161
			21.	Control Files Stored in the .SpaceMan Directory	161

About This Book

This manual provides information to help you perform the following tasks on your workstation:

- Install the HSM client
- Register your workstation
- Select your options for space management
- Select space management options
- Migrate your files automatically or manually
- Back up or archive your migrated or premigrated files
- Recreate your stub files
- Recall your migrated files
- Reconcile your file systems
- Schedule services.

The terms *hierarchical storage management (HSM)* and *space management* have the same meaning throughout this manual.

The HSM client includes two installed versions: a graphical user interface (GUI) client, and a command-line client. This manual provides information for both.

You must install the IBM Tivoli Storage Manager backup-archive client on your workstation before you can use the HSM client. This manual does not provide information about using the backup-archive clients. IBM Tivoli Storage Manager is a separate client-server licensed program that provides storage management services.

In addition to the HSM client, the following programs are included on several operating systems:

- A backup-archive client program that backs up and archives files from your workstation or file server to storage, and restores and retrieves backup versions and archived copies of files to your local file systems.
- An application program interface (API) that you can use to enhance an existing application with storage management services. When you register an application with a server as a client node, you can back up, restore, archive, and retrieve objects from storage.
- A Web backup-archive client that an authorized administrator, support person, or end user can use to perform backup, restore, archive, and retrieve services using a Web browser on a remote machine.
- An administrative client program that you can access from a Web browser or from the command line. An administrator controls and monitors server activities, defines storage management policies for backup, archive, and space management services, and sets up schedules to perform these services at regular intervals.
- A server program in a UNIX environment that performs either as a backup and archive server, or as a migration server for distributed workstations and file servers. The server program also supplies HSM services.

Who Should Read This Manual

This manual includes instructions to set up and maintain the HSM client on a workstation. You should be familiar with your workstation and your operating system. As root user, you should be familiar with your system administration. In addition, you should be familiar with the Tivoli Storage Manager backup-archive client.

IBM Tivoli Storage Manager Web site

Technical support information and publications are available at the following address:

<http://www.ibm.com/software/sysmgmt/products/support/IBMTivoliStorageManager.html>

By accessing the Tivoli Storage Manager home page, you can access subjects that interest you. You can also access current Tivoli Storage Manager product information.

The IBM Tivoli Storage Manager home page **Self help** section provides customers with a knowledge base of articles and information on issues with IBM Tivoli Storage Manager that they might be experiencing.

Conventions Used in This Manual

This manual uses the following typographical conventions:

Example	Description
<code>autoexec.ncf</code>	A series of lowercase letters with an extension indicates program file names.
archive	Boldface type indicates a command that you type on a command line.
<i>dateformat</i>	Boldface italic type indicates an option.
<i>filespec</i>	Italicized type indicates either the name of a parameter, a new term, or a placeholder for information that you provide. Italics also are used for emphasis in the text.
<code>maxcmdretries</code>	Monospace type indicates fragments of a program or information as it might appear on a display screen.
plus sign (+)	A plus sign between two keys indicates that you press both keys at the same time.

Reading Syntax Diagrams

This section describes how to read the syntax diagrams that are used in this manual. To read a syntax diagram, follow the path of the line. Read from left to right, and top to bottom.

- The ►— symbol indicates the beginning of a syntax diagram.
- The —> symbol at the end of a line indicates that the syntax diagram continues on the next line.
- The ►— symbol at the beginning of a line indicates that a syntax diagram continues from the previous line.
- The —<◄ symbol indicates the end of a syntax diagram.

Syntax items, such as a keyword or a variable, can be:

- On the line (required element)
- Above the line (default element)
- Below the line (optional element).

Table 1 describes the syntax diagram items and provides examples.

Table 1. Description of Syntax Diagram Items

Syntax Diagram Description	Example
<p>Abbreviations:</p> <p>Uppercase letters indicate the shortest acceptable form. If an item appears entirely in uppercase letters, you cannot shorten it.</p> <p>Type the item in any combination of uppercase or lowercase letters. In this example, you can enter any of these choices either in uppercase or lowercase letters: KEYWO, KEYWORD, or KEYWOrd.</p>	
<p>Symbols:</p> <p>Enter these symbols <i>exactly</i> as they appear in the syntax diagram.</p>	<ul style="list-style-type: none"> * Asterisk { } Braces : , = Equal sign - Hyphen () Parentheses . Space
<p>Variables:</p> <p>Italicized lowercase items (<i>var_name</i>) indicate variables.</p> <p>In this example, you can specify a <i>var_name</i> when you enter the KEYWOrd command.</p>	
<p>Repetition:</p> <p>An arrow returning to the left indicates that you can repeat the item.</p> <p>A character or space <i>within</i> the arrow indicates that you must separate repeated items with a character or space.</p> <p>A footnote indicates the number of times that you can repeat the item.</p>	  
<p>Notes:</p>	<p>1 Specify <i>repeat</i> as many as five times.</p>

Table 1. Description of Syntax Diagram Items (continued)

Syntax Diagram Description	Example
<p>Required Choices:</p> <p>When two or more items are in a stack and one of them is on the line, you <i>must</i> select one item.</p> <p>In this example, you <i>must</i> select either A, B, or C.</p>	
<p>Optional Choice:</p> <p>When an item is <i>below</i> the line, that item is optional. In the first example, you can select either A or nothing at all.</p> <p>When two or more items are in a stack below the line, all of them are optional. In the second example, you can select A, B, C, or nothing at all.</p>	
<p>Defaults:</p> <p>Defaults are above the line. The default is selected unless you override it, or you can select the default explicitly. To override the default, include an option from the stack below the line.</p> <p>In this example, A is the default. Select either B or C to override A.</p>	
<p>Repeatable Choices:</p> <p>A stack of items followed by an arrow returning to the left indicates that you can select more than one item, or in some cases, repeat a single item.</p> <p>In this example, you can select any combination of A, B, or C.</p>	
<p>Syntax Fragments:</p> <p>Some diagrams, because of their length, must fragment the syntax. The fragment name displays between vertical bars in the diagram. The expanded fragment displays between vertical bars in the diagram after a heading with the same fragment name.</p>	<p>▶▶ The fragment name ▶▶</p> <p>The fragment name:</p>

Related Information

Table 2 contains a list of the manuals that are a part of the IBM Tivoli Storage Manager client library. Some of these manuals might be referred to in this manual.

Order manuals through your Tivoli representative or the Tivoli branch office that serves your locality.

Table 2. Tivoli Storage Manager Client Manuals

Publication Title	Order Number
<i>IBM Tivoli Storage Manager for Macintosh: Backup-Archive Client Installation and User's Guide</i>	GC32-0787
<i>IBM Tivoli Storage Manager Messages</i>	GC32-0767
<i>IBM Tivoli Storage Manager for NetWare: Backup-Archive Client Installation and User's Guide</i>	GC32-0786
<i>IBM Tivoli Storage Manager for UNIX: Backup-Archive Clients Installation and User's Guide</i>	GC32-0789
<i>IBM Tivoli Storage Manager Using the Application Program Interface</i>	GC32-0793
<i>IBM Tivoli Storage Manager for Windows: Backup-Archive Clients Installation and User's Guide</i>	GC32-0788
<i>IBM Tivoli Storage Manager for AIX Storage Agent User's Guide</i>	GC32-0771
<i>IBM Tivoli Storage Manager Publications CD</i>	SK3T-8176

Downloading or ordering publications

All Tivoli publications are available for electronic download or order from the IBM Publications Center: <http://www.ibm.com/shop/publications/order/>.

If you have questions or comments regarding Tivoli publications and product documentation, please visit <http://www.ibm.com/software/tivoli/contact.html> to send an e-mail.

The International Technical Support Center (ITSC) publishes Redbooks, which are books on specialized topics such as using Tivoli Storage Manager to back up databases. You can order publications through your IBM representative or the IBM branch office serving your locality. You can also search for and order books of interest to you at the IBM Redbooks Web site at this address:

<http://www.ibm.com/redbooks/>

Tivoli Field Guides are designed to address specific technical scenarios or concepts that are often complex to implement or difficult to understand. All completed field guides are available free to registered customers and internal IBM employees at the following Web site:

http://www.ibm.com/software/sysmgmt/products/support/Field_Guides.html

Contacting customer support

For support for this or any Tivoli product, you can contact Tivoli Customer Support in one of the following ways:

- Visit the IBM Tivoli Storage Manager technical support Web site at:
<http://www.ibm.com/software/sysmgmt/products/support/IBMTivoliStorageManager.html>

- Submit a problem management record (PMR) electronically at **IBMSERV/IBMLINK**. You can access the IBMLINK from the IBM Web site at: <http://www.ibm.com/ibmlink/>
- Submit a problem management record (PMR) electronically from the IBM Web site at: <http://www.ibm.com/software/support/probsub.html>.

Customers in the United States can also call 1-800-IBM-SERV (1-800-426-7378).

International customers should consult the Web site for customer support telephone numbers.

You can also review the *IBM Software Support Guide*, which is available on our Web site at <http://techsupport.services.ibm.com/guides/handbook.html>.

When you contact IBM Software Support, be prepared to provide identification information for your company so that support personnel can readily assist you. Company identification information is needed to register for online support available on the Web site.

The support Web site offers extensive information, including a guide to support services (IBM Software Support Guide); frequently asked questions (FAQs); and documentation for all IBM Software products, including Release Notes, Redbooks, and white papers, defects (APARs), and solutions. The documentation for some product releases is available in both PDF and HTML formats. Translated documents are also available for some product releases.

We are very interested in hearing about your experience with Tivoli products and documentation. We also welcome your suggestions for improvements. If you have comments or suggestions about our documentation, please complete our customer feedback survey at:

<http://www.ibm.com/software/sysmgmt/products/support/IBMTivoliStorageManager.html>

by selecting the Feedback link in the left navigation bar.

Reporting a problem

Please have the following information ready when you report a problem:

- The IBM Tivoli Storage Manager server version, release, modification, and service level number. You can get this information by entering the **query status** command at the IBM Tivoli Storage Manager command line.
- It is recommended that you use the IBM Tivoli Storage Manager client **query systeminfo** command with the **filename** option to gather IBM Tivoli Storage Manager system information and output this information to a file. This information is intended primarily as an aid for IBM support to assist in diagnosing problems.
- The IBM Tivoli Storage Manager client version, release, modification, and service level number. You can get this information by entering `dsmc` at the command line.
- The communication protocol (for example, TCP/IP), version, and release number you are using.
- The activity you were doing when the problem occurred, listing the steps you followed before the problem occurred.
- The exact text of any error messages.

Internet

You can get additional information through an anonymous FTP server, [ftp://ftp.software.ibm.com](ftp://ftp.software.ibm.com/storage/tivoli-storage-management). IBM Tivoli Storage Manager information is in the `/storage/tivoli-storage-management` directory.

A newsgroup, listserv@marist.edu, is implemented by a third party. IBM supports this newsgroup on a best-effort basis only. To participate in user discussions of IBM Tivoli Storage Manager you can subscribe to the ADSM-L list server. This is a user forum maintained by Marist College. While not officially supported by IBM, IBM Tivoli Storage Manager developers and other IBM support staff also participate on an informal, best-effort basis. Because this is not an official IBM support channel, you should contact IBM Technical Support if you require a response specifically from IBM. Otherwise there is no guarantee that IBM will respond to your question on the list server.

You can subscribe by sending a note to the following e-mail address:

listserv@vm.marist.edu

The body of the message must contain the following:

SUBSCRIBE ADSM-L yourfirstname yourlastname

The list server will send you a response asking you to confirm the subscription request. Once you confirm your subscription request, the list server will send you further instructions. You will then be able to post messages to the list server by sending e-mail to:

ADSM-L@vm.marist.edu

If at a later time you want to unsubscribe from ADSM-L, you can send a note to the following e-mail address:

listserv@vm.marist.edu

The body of the message must contain the following:

SIGNOFF ADSM-L

You can also read and search the ADSM-L archives, join discussion forums, and access other resources at the following URL:

<http://www.adsm.org>

Chapter 1. Summary of the HSM Client

The IBM Tivoli Storage Manager for Space Management for UNIX (HSM) client migrates files from your local file system to storage and recalls them either automatically or selectively. Migrating files to a distributed storage device 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.

Your Tivoli Storage Manager administrator defines management classes to files. You, as root user:

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

These options and settings 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. Stub files that contain the necessary information to recall your migrated files remain on your local file system so that the files appear to reside locally. When you access migrated files, they are recalled automatically to your local file system. Contrast this with archiving, which completely removes files from your local file system.

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

File migration, unlike file backup, does not protect against accidental file deletion, file corruption, or disk failure. Continue to back up your files whether they reside on your local file system or in storage. You can use the IBM Tivoli Storage Manager backup-archive client to back up and restore migrated files in the same manner as you would back up and restore files that reside on your local file system. If you accidentally delete stub files from your local file system, or if you lose your local file system, you can restore the stub files.

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

Migrating Files from Your Local File System

Once file migration begins, the HSM client sends a copy of your file to a server and replaces the original file with a stub file on your local file system. The file appears as though it resides on your local file system. A *stub file* is a small file that contains required information to locate and recall a migrated file, and to respond to specific UNIX commands without recalling the file. The server places your migrated files in storage volumes on disk devices, or devices that support removable media, such as tape. Your Tivoli Storage Manager administrator defines and groups the storage volumes into storage pools.

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

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

Demand migration responds to an out-of-space condition on your local file system. Demand migration starts automatically if your file system runs out of space. As files are migrated, space becomes available on your file system and the process that caused the out-of-space condition continues.

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

Premigrating Files from Your Local File System

For faster migration, use the premigration process to prepare your files for automatic migration. 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 HSM client quickly changes premigrated files to migrated files without requiring additional time to copy them to storage. The HSM client verifies that files have not changed since they were premigrated. Copies of the files are replaced with stub files on your local file system.

Files are premigrated each time automatic migration completes if:

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

Recalling Your Migrated Files

You can recall a migrated file to your local file system from storage using either transparent recall or selective recall.

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, or for a specific execution of a command (AIX JFS only), you change how the HSM client recalls a migrated file.

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

Migrate-on-close

(Valid for AIX JFS file systems only)

Temporarily recalls a migrated file to its originating file system. The recalled file remains on your local file system only for as long as it is open. When you close the unmodified file, a stub file replaces it on your local file

system. The file is migrated again. It is not necessary to send a copy of the file to storage because you did not modify the file. The copy that currently resides in storage remains valid.

Normal

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.

Read-without-recall

(Valid for AIX JFS 4.3 file systems only)

Reads a migrated file from storage without storing it on your local file system. The HSM client reads information sequentially from the migrated file, and caches information that it reads from the file into a memory buffer on your workstation. The read-without-recall mode is intended for single-access, sequential reads of non-executable files.

Partial file recall

(Valid for AIX GPFS HSM only)

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. When HSM intercepts a read request for a file configured for partial file recall, it will calculate which portion of the file to recall based on the offsets contained in the read request. This results in time and disk space savings, since only a portion of the file is recalled, using less local disk space. See “Recall Your Files in an AIX GPFS Environment” on page 74 for more information.

Streaming recall mode

(Valid for AIX GPFS, Linux GPFS, and Solaris VxFS)

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.

See “Streaming Recall Mode” on page 74 for more information.

Selective recall returns specific, 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.

Reconciling Your File Systems

You, as root user, set reconciliation to automatically reconcile file systems at intervals. The default interval is every 24 hours. Set the value for ***reconcileinterval*** larger than zero. For a larger number of file systems, increase this value to reduce the impact that the ***ds mreconcile*** command might have on system performance. This reconciles your local file system with the server that you contact for space management services. When you reconcile your file system, you also update other space management-related information.

When you modify the data of a migrated or premigrated file, or erase a migrated or premigrated file from your local file system, you retain an obsolete copy of the file in storage. During reconciliation, any obsolete copies of migrated or premigrated files are marked for expiration. The copies expire and are removed from the server after the elapsed expiration date.

Chapter 2. Installing the HSM Client

Attention: Install and set up the IBM Tivoli Storage Manager for UNIX Backup-Archive Client *before* you install the HSM client. Both clients share common code and they use the same options files, communication protocols, node registration, and server file spaces.

This chapter provides installation and setup tasks for the HSM client. You must have root user authority to install, set up, and use the HSM client on your workstation.

You can add space management to the following operating systems:

- General Parallel File System (GPFS) 1.5 on a PSSP cluster (RS6000/SP) or on an AIX cluster (IBM Regatta pSeries 690). Both cluster systems require AIX 5.1 and PSSP 3.4.
- GPFS 2.2 or higher on IBM @server Cluster 1350 with CSM 1.3 (Cluster Systems Management).
- Journaled File System (JFS) on an AIX 5.1 or an AIX 5.2 workstation
- Veritas File System (VxFS) on a Sun Solaris 7, 8, or 9 workstation
- Veritas File System (VxFS) 3.5, Veritas Volume Manager (VxVM) 3.5, and OnlineJFS 3.3 on a HP-UX 11i (11.11) workstation and server.

The HSM client must communicate with a Tivoli Storage Manager server.

Additional components for each operating system include a command-line help and a graphical user interface help.

Install the HSM client on your workstation and register it as a client node with a IBM Tivoli Storage Manager server. Assign this node to a policy domain containing one or more management classes that are configured to support space management services. When you complete the installation, continue with Chapter 3, “Selecting Your Options,” on page 25 to modify your `dsm.sys` file and your `dsm.opt` file for space management. Both files are located in the `/usr/tivoli/tsm/client/ba/bin` directory.

Attention: Before you *reinstall* the HSM client, stop all activity and do not access any files on file systems to which you added space management. The install process will fail otherwise.

Table 3 displays the page location for you to go to begin installation of the HSM client.

Table 3. HSM Client Installation Procedures

Installation Procedure	Page
Installing the HSM Client on AIX GPFS File Systems	6
Installing the HSM Client on AIX JFS File Systems	11
Installing the HSM Client on HP-UX VxFS File Systems	14
Installing the HSM Client on Linux86 GPFS File Systems	16
Installing the HSM Client on Solaris VxFS File Systems	19

Moving Your Files to Another Directory

On AIX, the backup-archive client and the common files are installed in the `/usr/tivoli/tsm/client/ba/bin` directory. The HSM client files are installed in the `/usr/tivoli/tsm/client/hsm/bin` directory.

On Solaris, HP-UX, and Linux86, the backup-archive client and the common files are installed in the `/opt/tivoli/tsm/client/ba/bin` directory. The HSM client files are installed in the `/opt/tivoli/tsm/client/hsm/bin` directory.

If you move HSM files to another directory, perform the following steps:

1. Verify that the permissions of the installed files did not change.
2. Update the symbolic links for the installed files in the `/usr/bin` directory, and in the directory containing symbolic links for each language package that you install; for example, `/usr/lib/nls/msg/en_US`.
3. Set the `DSM_DIR` environment variable to point to the newly-installed directory.

Installing the HSM Client on AIX GPFS File Systems

Attention: The HSM client for AIX GPFS and the HSM client for AIX JFS are exclusive of each other. During installation, the program checks for the existence of the other. If the program detects either client, installation stops.

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
```

The HSM client for AIX GPFS supports RS6000/SP systems (PSSP cluster) and IBM Regatta pSeries (AIX HACMP/ES cluster).

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

1. Define which node within your GPFS nodeset will become the HSM session node (this is the node on which the GPFS file system(s) are HSM managed and all HSM daemons are running on).
2. Define which node(s) within your GPFS nodeset will become HSM backup node(s) in case the HSM session node fails.
3. Define which GPFS file systems will be HSM managed.
4. Install Tivoli SpaceManager on the HSM session node (1). Make sure that after installation, the **dsmrecalld** daemon is up.
5. Unmount all GPFS file systems on all nodes within the GPFS nodeset, that will be HSM managed.
6. Activate DMAPI management for these GPFS file systems:

```
mmchfs -z yes <device>
```
7. Remount all GPFS file systems on all nodes within the GPFS nodeset.
8. Define your Tivoli Storage Manager server and client configuration in the `dsm.sys` and `dsm.opt` files.
9. Stop and restart the HSM daemons:
 - `kill -15 <pid of dsmonitord>` stops **dsmonitord**.
 - `kill -15 <pid of dsmscoutd>` stops **dsmscoutd**.
 - `'dmkilld'` stops and restarts the **dsmrecalld**.

Note: Because the failover technique is enabled by default after installation, the **dsmwatchd** will automatically restart the HSM daemons.
10. On the HSM session node, add HSM management to at least one GPFS file system (`dsmmigfs add <fs>`). The HSM daemons will detect the initial state and assign this node the instance number one.
11. Install Tivoli SpaceManager on the HSM backup node(s). If you plan to setup more than one HSM backup node, do the installation one after the other. The HSM backup node(s) will have subsequent instance numbers assigned.

Note: The HSM daemon **dsmwatchd** needs to be up and running on the HSM backup node(s). The other HSM daemons are started automatically if the HSM session node fails.

When you install the HSM client for GPFS on a PSSP or AIX cluster file system, the install process:

- Stops any space management daemons that are running.
- Removes any statement from the `/etc/inittab` file that loads the **dsmwatchd** command at system startup.
- Removes any statement from the `/var/mmfs/etc/gpfsready` script file that loads the other space management daemons at GPFS system startup.
- Extracts the HSM modules.
- Adds a statement to the `/etc/inittab` file that loads the **dsmwatchd** command at system startup.
- Adds a statement to the `/var/mmfs/etc/gpfsready` script file that loads the other space management daemons at GPFS system startup. If the file does not exist, the install process creates the following HSM-specific files or objects for each GPFS node set in the SDR for PSSP cluster:

SDR	Files
Files	dsm.opt.x dsm.sys.x dsmmigfstab.x
	Where <i>x</i> is the node-set ID
Class	DSMLock
Class Objects	dsm.opt.x dsm.sys.x dsmmigfstab.x dsmmigfs.x
	Where <i>x</i> is the node-set ID

Note: The HSM client must have read-write access to the SDR (for PSSP cluster systems).

If the file does not exist, the install process creates the following HSM-specific files or objects for each GPFS node set in the Tivoli Storage Manager client internal SDR for an AIX cluster:

- dsmmigfstab
- dsm.opt
- dsm.sys
- HSMnode.xxx for each cluster node the **dsmwatchd** is active where *xxx* is a unique instance number within the given GPFS node set.
- Starts the space management daemons.

The sections that follow describe client environment requirements and installation instructions for the HSM client. For current installation and configuration information for the HSM client, refer to the readme file that is shipped on the product installation media.

Environment Requirements

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

Hardware Requirements

- PSSP cluster (RS6000/SP) with GPFS 1.5 or AIX cluster (IBM Regatta pSeries 690) with PSSP 3.4 and GPFS 1.5
- Disk space: See the readme file that is shipped on the product installation media
- Memory: 64 MB

Software Requirements

- Common Desktop Environment (CDE)
- GPFS 1.5 or 2.1 PTF U488745
- Motif 1.2 or 2.0
- X Window System X11R6

Operating System Requirements

- AIX 5.1
- AIX Windows Environment/6000 for AIX 5.1 for the administrative client

Communication Methods

This table displays what you must install to communicate with your server.

To use this communication method	To connect to these servers	Install this on your workstation
TCP/IP	AIX, HP-UX, Linux, OS/390, OS/400 PASE, Solaris, VM, Windows, z/OS	TCP/IP (Standard with supported AIX/6000 platforms)
Shared memory	AIX	TCP/IP (Standard with supported AIX platforms)

Installation

Attention: The GPFS-enhanced Tivoli Storage Manager backup-archive client (tivoli.tsm.client.ba.gpfs.aix51.32bit) is a prerequisite for the HSM client on AIX GPFS file systems. You must install this file set. The backup-archive client is enhanced to handle any file systems that the HSM client handles for AIX GPFS.

Before you install the HSM client, verify that there are no HSM activities processing on your managed file systems.

The following packages are available on the installation media in the /usr/sys/inst.images directory:

Package	Installs	Into
tivoli.tsm.client.ba.gpfs.aix51.32bit	The backup-archive client for AIX 5.1 32bit and AIX 5.1 32bit	The /usr/tivoli/tsm/client/ba/bin directory
tivoli.tsm.client.hsm.gpfs.aix51.32bit	The HSM client for AIX 5.1 32bit	The /usr/tivoli/tsm/client/hsm/bin directory

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

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

Install the HSM Client from the Command Line

If you are not installing the HSM client from a CD-ROM, enter the following command from the AIX command line:

```
/usr/sbin/inutoc <dir>
```

Where <dir> is the directory in which the HSM image resides. A .toc file is created in this directory. Continue with step 2 located in “Install the HSM Client from the CD-ROM.”

Install the HSM Client from the CD-ROM

If you are installing the HSM client from a CD-ROM, insert the CD-ROM into the drive. Then, follow the steps below.

1. Mount your CD-ROM drive.
2. Type **smitty install** and press **Enter**.

3. Select **Install and Update Software** and press **Enter**.
4. Select **Install and Update From ALL Available Software** and press **Enter**.
5. At the **INPUT device/directory for software** prompt, press **F4**.
6. Select the CD-ROM device that contains the installation CD-ROM (for example, /dev/cd0), or specify the directory that contains the installation images.
7. Press **Enter**.
8. At the **SOFTWARE to install** prompt, press **F4**.
9. Scroll through the file sets and press **F7** to select each Tivoli Storage Manager file set that you want to install.
10. Press **Enter**. A minimum installation of HSM requires the following:
 - Tivoli Storage Manager backup-archive client common files
 - Tivoli Storage Manager backup-archive client options
 - HSM client options.
11. Select any options that you want and press **Enter** to begin the installation. For more information about options, see Chapter 3, “Selecting Your Options,” on page 25 or Chapter 12, “Using Options,” on page 89.
12. Unmount the CD-ROM drive.

Continue with “**Registering Your Workstation with a Server**” on page 21.

Installing the HSM Client on AIX JFS File Systems

Attention: The HSM client for AIX JFS and the HSM client for AIX GPFS are exclusive of each other. During installation, the program checks for the existence of the other. If the program detects either client, installation stops.

When you install the HSM client on an AIX JFS file system, the install process:

- Unmounts the file system migrator (fsm) from all file systems with a file system type of fsm. These file systems exist only if you previously installed the HSM client *and* you added space management.
- Stops any space management daemons that are running.
- Unloads the HSM kernel extension, if it exists.
- Removes any statement from the `/etc/inittab` file that loads the kernel extension.
- Extracts the HSM modules.
- Loads the HSM kernel extension.
- Mounts the fsm.
- Starts the space management daemons.
- Adds a statement to the `/etc/inittab` file that loads the kernel extension (kext) at system startup.

See Appendix A, “Using the Installfsm Program,” on page 159 for more information about the installfsm program.

The sections that follow describe client environment requirements and installation instructions for the HSM client. For current installation and configuration information for the HSM client, refer to the readme file that is shipped on the product installation media.

Environment Requirements

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

Hardware Requirements

- RS/6000
- Disk space: See the readme file that is shipped on the product installation media
- Memory: 64 MB

Software Requirements

- Common Desktop Environment (CDE)
- Motif 1.2 or 2.0
- X Window System X11R6

Operating System Requirements

- AIX 5.1 or AIX 5.2
- X11/Motif/CDE for the administrative client

Communication Methods

This table displays what you must install to communicate with your Tivoli Storage Manager server.

To use this communication method	To connect to these Tivoli Storage Manager servers	Install this on your workstation
TCP/IP	AIX, HP-UX, Linux, OS/390, OS/400 PASE, Solaris, VM, Windows, z/OS	TCP/IP (Standard with supported AIX/6000 platforms)
Shared memory	AIX	TCP/IP (Standard with supported AIX platforms)

Installation

Attention: Install and set up the Tivoli Storage Manager for UNIX Backup-Archive Client before you install the HSM client. Some common files, such as `dsm.sys` and `dsm.opt`, are installed in the `/usr/tivoli/tsm/client/ba/bin` directory.

The following packages are available on the installation media in the `/usr/sys/inst.images` directory:

Package	Installs	Into
<code>tivoli.tsm.client.ba.aix32bit</code>	The backup-archive client for AIX JFS 32bit	The <code>/usr/tivoli/tsm/client/ba/bin</code> directory
<code>tivoli.tsm.client.hsm.jfs.aix32bit</code>	The HSM client for AIX JFS 32bit	The <code>/usr/tivoli/tsm/client/hsm/bin</code> directory
<code>tivoli.tsm.client.ba.aix64bit</code>	The backup-archive client for AIX JFS 64bit	The <code>/usr/tivoli/tsm/client/ba/bin</code> directory
<code>tivoli.tsm.client.hsm.jfs.aix64bit</code>	The HSM client for AIX JFS 64bit	The <code>/usr/tivoli/tsm/client/hsm/bin</code> directory

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

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

Install the HSM Client from the Command Line

If you are not installing the HSM client from a CD-ROM, enter the following command from the AIX command line:

```
/usr/sbin/inutoc <dir>
```

Where `<dir>` is the directory in which the HSM image resides. A `.toc` file is created in this directory. Continue with step 2 located in “Install the HSM Client from the CD-ROM.”

Install the HSM Client from the CD-ROM

If you are installing the HSM client from a CD-ROM, insert the CD-ROM into the drive and continue with the steps below.

1. Mount the CD-ROM drive.
2. Type **smitty install** and press **Enter**.
3. Select **Install and Update Software** and press **Enter**.

4. Select **Install and Update From ALL Available Software** and press **Enter**.
5. At the **INPUT device/directory for software** prompt, press **F4**.
6. Select the CD-ROM device that contains the installation CD-ROM (for example, /dev/cd0), or specify the directory that contains the installation images.
7. Press **Enter**.
8. At the **SOFTWARE to install** prompt, press **F4**.
9. Scroll through the file sets and press **F7** to select each Tivoli Storage Manager file set that you want to install.
10. Press **Enter**. A minimum installation of the HSM client requires the following:
 - Tivoli Storage Manager client common files
 - Tivoli Storage Manager backup-archive client options
 - HSM client options.
11. Select any options that you want and press **Enter** to begin the installation. For more information about options, see Chapter 3, “Selecting Your Options,” on page 25 or Chapter 12, “Using Options,” on page 89.
12. Unmount the CD-ROM drive.

Obtain VFS and Group Numbers

When you install HSM on an AIX JFS file system, the HSM client obtains a virtual file system (vfs) number. Because numbers zero through seven are reserved, the program attempts to obtain a vfs number in the range of eight through fifteen. The HSM client begins with the number fifteen and descends in that range to locate an unused number. If a number is not available, an error message displays. You must:

1. Check the /etc/vfs file for the current use of vfs numbers.
2. Release a number in the range of eight through fifteen.
3. Start Tivoli Storage Manager installation again.

The HSM client also obtains group numbers for HSM execution modes. The program selects an unused range of group numbers beginning with the largest possible group number and descending to the smallest possible group number. The range is xxxxxx00 through xxxxxx7F.

Continue with “Registering Your Workstation with a Server” on page 21.

Installing the HSM Client on HP-UX VxFS File Systems

The sections that follow describe client environment requirements and installation instructions for the HSM client on HP-UX file systems. For current installation and configuration information for the HSM client, refer to the readme file that is shipped on the product installation media.

Environment Requirements

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

Hardware Requirements

- HP 9000 Series 700 or 800 workstation or server with PA-RISC CPU
- Disk space: See the readme file that is shipped on the product installation media
- Memory: 256 MB

Software Requirements

- Common Desktop Environment (CDE)
- Motif 2.0
- Veritas file systems (VxFs) 3.5 with performance Patch from Vertias (see File README for more information), Veritas Volume Manager (VxVM) 3.5, OnlineJFS (3.3) or (3.5)
- X Window System X11R6

Operating System Requirements

- HP-UX 11.i (11.11)

Communication Methods

This table displays what you must install to communicate with your Tivoli Storage Manager server.

To use this communication method	To connect to these Tivoli Storage Manager servers	Install this on your workstation
TCP/IP	AIX, HP-UX, Linux, OS/390, OS/400 PASE, Solaris, VM, Windows, z/OS	TCP/IP (Standard with HP-UX)
Shared memory	HP-UX	TCP/IP (Standard with HP-UX)

Installation

Attention: Install and set up the IBM Tivoli Storage Manager for UNIX Backup-Archive Client *before* you install the HSM client. Install HSM on HP-UX VxFS file systems *only*. Do not attempt to install HSM on other HP-UX file systems.

The Veritas file system requires that the kernel extension (thsmdmapi) is loaded before you can use the HSM client. This kernel extension loads automatically during startup.

To install from the CD-ROM, log in as the root user, mount the CD-ROM to /cdrom, and change directory to tsmcli/hp11. If you downloaded from ftp, go to the directory where the installable image is located. Enter the following command to install the IBM Tivoli Storage Manager for UNIX Backup-Archive Client:

```
| /usr/sbin/swinstall -x mount_all_filesystems=false -v -s 'pwd'/  
| TIVsmC TIVsm
```

| **Enter the following command to install the HSM client:**

```
| /usr/sbin/swinstall -x mount_all_filesystems=false -v -s 'pwd'/ TIVsmChsm  
| TIVsm.CLIENT_HSM
```

| **Note:** 'pwd' may be used instead of the absolute name of the current directory.

| **Continue with “Registering Your Workstation with a Server” on page 21.**

Installing the HSM Client on Linux86 GPFS File Systems

For an initial installation on Linux86 cluster systems (IBM @server Cluster 1350), install the HSM client on only one node and enable at least one GPFS file system for HSM on that node (**dsmmigfs add**). The HSM daemons will detect the initial state and assign this node the instance number one. The first HSM-managed file system must be defined on this node. Then, select other nodes if you want to install the HSM client for failover, or to manage other GPFS file systems.

Install Tivoli Storage Manager on all cluster nodes that are part of a GPFS node set and that will become a DMAPI session node in the event of failover.

When you install the HSM client for GPFS on a Linux86 cluster file system, the install process:

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

The sections that follow describe client environment requirements and installation instructions for the HSM client. For current installation and configuration information for the HSM client, refer to the readme file that is shipped on the product installation media.

Environment Requirements

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

Hardware Requirements

- IBM @server Cluster 1350 with CSM 1.3 (Cluster Systems Management) and GPFS 2.2 or higher. See IBM Redbook *Linux Clustering with CMS and GPFS*, SG24-6601 for more information.
- Disk space: See the readme file that is shipped on the product installation media
- Memory: 64 MB

Software Requirements

- GPFS 2.2 or higher
- Linux kernel 2.4.4 or higher
- glibc 2.2.2 or higher
- libstdc++2.9.6 or higher
- X Window System X11R6 (for end user GUI only)
- RPM 3.0.0 or higher, 4.0
- Java JRE 1.3.1 or higher for the Web client and Java GUI
- A Netscape 4.7 or higher browser or Microsoft Internet Explorer 5.0 or higher browser

Operating System Requirements

- SuSE Linux Enterprise Server 8.0
- GPFS 2.2

Communication Methods

This table displays what you must install to communicate with your server.

To use this
communication
method

To connect to these servers

Install this on your workstation

TCP/IP

AIX, HP-UX, Linux, OS/390,
OS/400 PASE, Solaris, VM,
Windows, z/OS

TCP/IP (Standard with Linux)

Installation

Attention: The GPFS-enhanced Tivoli Storage Manager backup-archive client (TIVsm-BA.i386.rpm) is a prerequisite for the HSM client on Linux86 GPFS file systems. You must install this file set. The backup-archive client is enhanced to handle any file systems that the HSM client handles for Linux86 GPFS.

Before you install the HSM client, verify that there are no HSM activities processing on your managed file systems.

The following packages are available on the installation media in the `/opt/sys/inst.images` directory:

Package	Installs	Into
TIVsm-BA.i386.rpm	The Tivoli Storage Manager backup-archive client (command-line and native GUI), the administrative client (command-line), the Web backup-archive client, and the Tivoli Storage Manager documentation.	The <code>/opt/tivoli/tsm/client/ba/bin</code> directory
TIVsm-HSM.i386.rpm	The HSM client for Linux86	The <code>/opt/tivoli/tsm/client/hsm/bin</code> directory

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

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

If you are installing the HSM client from a CD-ROM, insert the CD-ROM into the drive. Then, follow the steps below.

1. Log in as the root user and mount the CD-ROM to the `/cdrom` directory.
2. Enter the following directory path where the installation packages reside on the CD:

`/cdrom/tsmcli/linux86`

3. Install the Tivoli Storage Manager clients in the order that is presented in Table 4. During installation, these packages are installed in unique directories.
4. Respond **Yes (y)** to the questions.
5. Restart your system when you complete the installation.

Table 4. Package Names and Descriptions

To Install This	Enter This Command
API	rpm -i TIVsm-API.i386.rpm Installs the API files in the /opt/tivoli/tsm/client/api/bin directory.
Backup-archive client (CLI and GUI) Web client	rpm -i TIVsm-BA.i386.rpm To circumvent the dependence check, use the --nodeps option. Please be aware to check the dependencies manually:
Administrative client	First: rpm -i --nodeps TIVsm-API.i386.rpm Then: rpm -i --nodeps TIVsm-BA.i386.rpm
Documentation	Installs the backup-archive client (command-line and GUI) and the Web client in the /opt/tivoli/tsm/client/ba/bin directory. Installs the the administrative client (command-line) in the /opt/tivoli/tsm/client/admin/bin directory. Installs the Tivoli Storage Manager documentation files in the /opt/tivoli/tsm/client/books directory.
HSM client	rpm -i TIVsm-HSM.i386.rpm Installs the HSM client files in the /opt/tivoli/tsm/client/hsm/bin directory.

Continue with “Registering Your Workstation with a Server” on page 21.

Installing the HSM Client on Solaris VxFS File Systems

Note: Install HSM on Solaris VxFS file systems *only*. Do not attempt to install HSM on other Solaris file systems.

The sections that follow describe client environment requirements and installation instructions for the HSM client on Solaris file systems. For current installation and configuration information for the HSM client, refer to the readme file that is shipped on the product installation media.

If you want to display the Tivoli Storage Manager help browser menus in your local language, the NLSPATH environment variable in the `/etc/profile` file must contain the following path:

```
/usr/dt/lib/nls/msg/%L/%N.cat
```

To enable the CDE help browser to use non-English fonts and messages, modify the NLSPATH environment variable:

```
NLSPATH=/usr/dt/lib/nls/msg/%L/%N.cat:$NLSPATH export NLSPATH
```

Environment Requirements

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

Hardware Requirements

- A SPARCstation or compatible workstation
- Disk space: See the readme file that is shipped on the product installation media
- Memory: 64 MB

Software Requirements

- Common Desktop Environment (CDE)
- Motif 2.0
- Veritas VxFS 3.3.2, or Veritas VxFS 3.4 patch level 2
- X Window System X11R6

Operating System Requirements

- Solaris 7, 8, or 9

Communication Methods

This table displays what you must install to communicate with your Tivoli Storage Manager server.

To use this communication method	To connect to these Tivoli Storage Manager servers	Install this on your workstation
TCP/IP	AIX, HP-UX, Linux, OS/390, OS/400 PASE, Solaris, VM, Windows, z/OS	Standard with Solaris
Shared memory	Solaris	TCP/IP (Standard with Solaris)

Installation

The Solaris file system requires that the kernel extension (adsmhsm) is loaded before you can use the HSM client. This kernel extension loads automatically during startup. Once the kernel extension is loaded, you cannot unload it. If you plan to install a new version of the HSM client, remove the HSM client and restart your system before you install a new version. If you do not plan to reinstall the HSM client, remove the HSM client and restart your system to unload the kernel extension.

Follow these steps to install the HSM client:

1. Log in as the root user and mount the CD-ROM to the /cdrom directory.
2. Install the Tivoli Storage Manager clients in the order that is presented in Table 5. During installation, these packages are installed in unique directories.
3. Respond **Yes (y)** to the questions.
4. Restart your system when you complete the installation.

Table 5. Package Names and Descriptions

To Install This	Enter This Command
API	<pre>pkgadd -d /cdrom/tsmcli/solaris/TIVsmCapi.pkg TIVsmCapi</pre> <p>Installs the API files in the /opt/tivoli/tsm/client/api/bin directory.</p>
Backup-archive client (CLI and GUI) Web client	<pre>pkgadd -d /cdrom/tsmcli/solaris/TIVsmCba.pkg TIVsmCba</pre> <p>Installs the backup-archive client and web client files in the /opt/tivoli/tsm/client/ba/bin directory.</p>
Administrative client	<p>Installs the administrative client in the /opt/tivoli/tsm/client/admin/bin directory.</p>
HSM client	<pre>pkgadd -d /cdrom/tsmcli/solaris/TIVsmChsm.pkg TIVsmChsm</pre> <p>Installs the HSM client files in the /opt/tivoli/tsm/client/hsm/bin directory.</p>
Client documentation	<pre>pkgadd -d /cdrom/tsmcli/solaris/TIVsmCdoc.pkg TIVsmCdoc</pre> <p>Installs the documentation files in the /opt/tivoli/tsm/client/books directory.</p>

Continue with “Registering Your Workstation with a Server” on page 21.

Registering Your Workstation with a Server

Before you request services from a Tivoli Storage Manager server, register your workstation as a single client node with that server. If you plan to use a web client with an enterprise management server, ensure that you are assigned an administrative user ID with system privilege, policy privilege, client access, or client-owner authority. Your Tivoli Storage Manager administrator assigns registration as either closed or open.

Use Closed Registration

With closed registration, your Tivoli Storage Manager administrator must register your workstation as a client node with the server. If your enterprise uses closed registration, provide the following information to your administrator:

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

Your administrator defines:

- The policy domain to which your client node belongs. 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.

Continue with “Starting the HSM Client” on page 23.

Use 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.

To start a session with the graphical user interface, enter the **dsm** command. To start a session with the command-line interface, enter the **dsmc** command.

The first time you start a session, 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.

To register your workstation with additional servers, enter the **dsm** command or the **dsmc** command with the **servername** option for each server. For example:

```
dsm -servername=dsm serv
```

Where **dsm serv** is the name of a server that you identified in your **dsm.sys** file. You are prompted for information to register your workstation with the server that you specify.

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.

Continue with “Starting the HSM Client” on page 23.

Changing your Password

Note: This task is optional.

If you installed the HSM client on your client node and you require a Tivoli Storage Manager password, set the **passwordaccess** option to *generate* in your `dsm.sys` file. Set the **mailprog** option in your `dsm.sys` file to send you the password each time it generates a new one. Tivoli Storage Manager encrypts and stores your password locally, and automatically generates a new password for your client node each time it expires. You are not prompted for a Tivoli Storage Manager password.

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

If you specify a migration server on 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 on either the **migrateserver** option or the **defaultserver** option, the password that you set applies to the named server in the first stanza of your `dsm.sys` file.

To change your password from the GUI:

1. Click **Utilities** on the menu bar.
2. Click **Change Password**. The Change Password window displays.
3. Enter your current Tivoli Storage Manager password in the Current password field.
4. Enter a new password in the New password field.
5. Enter the new password again in the Re-enter new password field.
6. Click **OK**. The Tivoli Storage Manager client updates your password for the server that your client node contacts for space management services.

Note: To change your Tivoli Storage Manager password from the command line, enter:

```
dsmssetpw oldpw newpw
```

Replace `oldpw newpw` with your old and new passwords.

Starting the HSM Client

Attention: Before you can start the graphical user interface (GUI), CDE must be installed and the X Window System must be running. It is not necessary for CDE to be running when you start the graphical user interface. You can use hierarchical storage management commands rather than the GUI to perform most space management tasks. The GUI is only available for root users. It is not available for HSM on AIX GPFS or HP-UX VxFS.

To start the GUI and open the Tivoli Storage Manager window, double-click the icon. If the Tivoli Storage Manager icon does not display, enter **dsmhsm**. A window displays from which you can select a task.

Continue with Chapter 3, “Selecting Your Options,” on page 25.

Chapter 3. Selecting Your Options

Attention: During installation *on AIX* of the IBM Tivoli Storage Manager for UNIX Backup-Archive Clients program, sample options files (dsm.sys.smp, and dsm.opt.smp) are placed in the /usr/tivoli/tsm/client/ba/bin directory. During installation *on Solaris, HP-UX, and Linux86* of the IBM Tivoli Storage Manager for UNIX Backup-Archive Clients program, sample options files (dsm.sys.smp, and dsm.opt.smp) are placed in the /opt/tivoli/tsm/client/ba/bin directory. If you are installing both the backup-archive client and the HSM client at the same time, copy and rename the sample options files and modify them for your use for both clients. If you previously installed the backup-archive client and you set up your options files, modify them for space management.

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

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

This chapter defines the space management options that you can set in both your dsm.sys file and your dsm.opt file. As optional tasks, you can assign management classes to your files and modify your include-exclude file. For more information about these options, see Chapter 12, “Using Options,” on page 89.

After you select the options for space management, mount all HSM-managed file systems either automatically or manually each time you restart your system.

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

Setting Options for Space Management

To use the HSM client, an authorized user must set options for space management. These options determine such things as:

- The Tivoli Storage Manager server to which your files migrate or premigrate
- How often space usage is checked on your file systems
- How often your file systems are automatically reconciled
- How often candidates are searched for automatic migrations
- How many automatic migration processes for each file system can migrate files in parallel
- How many days must elapse after a migrated or premigrated file is deleted or modified on your local file system before the copy in Tivoli Storage Manager storage expires and is removed.

Display Option Information

You can display information about options either from the command line or from the graphical user interface (GUI). From the GUI:

1. Click **Utilities**→**Display Options**.

2. Scroll up or down to view descriptions of client-server options.
3. Click **Help** to display descriptions of the options.
4. Click **Close** to close the Display Options window.

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

```
dsmmigquery -options
```

Edit Your Options Files

To edit your options files from the Backup-Archive Client for UNIX graphical user interface, perform these steps:

1. Select the **Utilities** menu → **Setup Wizard**.
2. Check **Help me configure the TSM Backup Archive Client**.
3. Check **Update my options file**.
4. Follow the instructions on your screen.

The sections that follow provide descriptions of each space management option that you can use. For more information about these options, see Chapter 12, “Using Options,” on page 89.

Set Options in Your dsm.sys File

In addition to backup and archive options, the dsm.sys file:

- Specifies which server your client node contacts for space management services
- Identifies the communication options for each server
- Specifies scheduling and space management options.

Table 6 on page 28 provides a brief description of each space management option that you can set in your dsm.sys file. The values that you set for the HSM client affect automatic migration, reconciliation, and recall. For more information about these options, see Chapter 12, “Using Options,” on page 89.

For information about setting up space management for AIX in an HACMP environment, see “Set Up Space Management for AIX in an HACMP Environment” on page 36.

For AIX GPFS and Linux86 GPFS file systems only: To access your data on the server in the event of failover, ensure that all HSM nodes participating in the failover environment share the same node name in your dsm.sys file.

With the exception of *errorprog*, place all space management options *at the beginning* of your dsm.sys file *before* any server stanzas that you define.

1. Obtain the server information from your Tivoli Storage Manager administrator.
2. Edit your dsm.sys file to include the server to which you want to connect for space management services.
3. Assign a name to the server to which you want to contact for space management services. For each Servername entry, include a COMMmethod entry to specify the communication method to use for client and server communications.
4. Enter a value for each option and remove the leading asterisk (*). You can specify options for more than one server.

You can specify both a default server and a migration server in your dsm.sys file. If you do not specify a server name on either the **defaultserver** option or the **migrateserver** option, then, the server that you do specify in the first stanza of your dsm.sys file becomes your default server.

If you specify a migration server with the **migrateserver** option, it overrides the server that you specified with the **defaultserver** option.

Figure 1 displays an example of a server stanza containing options for a server that you want to contact.

```

DEFAULTServer          server1
MIGRATEServer         server2
CHECKThresholds       2
CANDIDATESInterval    12
MAXCANDprocs          5
RECOncileinterval    12
MAXRECOncileproc     5
MAXThresholdproc     5
MINMIGFILESize       8192
MIGFILEEXPIration    10
MINRECALLdaemons     5
MAXRECALLdaemons     15
CHECKFororphans       no
MAXMIGRators          1
KERNElmessages       no
OVERLAPRECALL        no
Servername          server1
  COMMmethod          TCPip
  TCPPort             1500
  TCPServeraddress    almvmd.almaden.ibm.com
  Passwordaccess      generate
  Mailprog            /usr/bin/xsend root
  Groups              system tsm
  Users               steiner chron wang nguyen
  Incl excl           /adm/tsm/backup.excl
  ERRORProg           /bin/cat

Servername          server2
  COMMmethod          SNA1u6.2
  PARTner1uname       raptor
  TPname              appcdel
  CPICM0dename        appc
  Passwordaccess      generate
  Mailprog            /usr/bin/xsend root
  Groups              system tsm
  Users               sullivan tang stewart
  Incl excl           /adm/tsm/migrate.excl
  ERRORProg           /bin/cat

```

Figure 1. Example of a *dsm.sys* File with HSM Options

Table 6. Setting Space Management Options in Your *dsm.sys* File

Option	Description	Default
<i>candidatesinterval</i>	Specifies how often the <i>dsmscoutd</i> daemon searches for migration candidates in file systems to which you added space management. The range of values is one through 9999.	1
<i>checkfororphans</i>	Specifies whether or not the <i>dsmreconcile</i> command checks for orphans.	No
<i>checkthresholds</i>	Determines how frequently the space monitor daemon checks space usage on the file systems to which you add space management.	Every 5 minutes
<i>compression</i>	Compresses files before you send them to the Tivoli Storage Manager server.	No
<i>defaultserver</i>	Specifies the name of the default Tivoli Storage Manager server to contact for space management services if you do not specify a server name on the <i>migrateserver</i> option. You cannot override the default server for space management services in your <i>dsm.opt</i> file.	Server name in the first stanza
<i>errorprog</i>	Specifies a program to which you want to send a message if a severe error occurs during space management processing.	None
<i>kernelmessages</i>	Valid for AIX JFS file systems only Displays HSM-related messages during processing. If you change the value, it will not take effect until you restart your system.	Yes
<i>maxcandprocs</i>	Specifies the number of slave scout daemons that can scan for migration candidates in parallel on a system. The range of values is two through 20.	5
<i>maxmigrators</i>	Specifies the number of parallel migration sessions that you can perform. The range of values is one through 20.	1
<i>maxrecalldaemons</i>	Specifies the maximum number of recall daemons that you can run at one time. The range of values is two through 99.	20
<i>maxreconcileproc</i>	Specifies the maximum number of automatic reconciliation processes that you can start at one time. The range of values is one through 99.	3
<i>maxthresholdproc</i>	Specifies the maximum number of automatic threshold migration processes that you can start at one time. The range of values is one through 99.	3
<i>migfileexpiration</i>	Specifies the number of days that copies of migrated or premigrated files remain on the Tivoli Storage Manager server after they are modified on, or erased from, your local file system. The range of values is zero through 9999.	7 days
<i>migrateserver</i>	Specifies the Tivoli Storage Manager server to which you want to migrate files from your client node. Specify only one migration server for each client node. You cannot override the migration server in your <i>dsm.opt</i> file.	Specify a server name in the first stanza
<i>minmigfilesize</i>	Specifies the minimum file size for migration candidates that is smaller or larger than the stub file size. The range of values is the fragment size of the file system through 2147483647.	Zero
<i>minrecalldaemons</i>	Specifies the minimum number of recall daemons that you want to run at the same time to perform recalls for your client node. The range of values is one through 99.	3

Table 6. Setting Space Management Options in Your *dsm.sys* File (continued)

Option	Description	Default
<i>overlaprecall</i>	Valid for AIX JFS file systems only Specifies whether migrated files must be recalled immediately even if a process is accessing only the resident part of the file. Specify Yes or No. This option is valid for read-only operations on a file.	No
<i>reconcileinterval</i>	Specifies how often automatic reconciliation of file systems occurs to which you add space management. The range of values is zero through 9999.	Every 24 hours

Set Options in Your *dsm.opt* File

The *dsm.opt* file contains options that affect the backup-archive client. However, some options, such as ***optionformat*** and ***restoremigstate***, affect both the backup-archive client and the HSM client. Table 7 provides a brief description of the space management options that you can set in your *dsm.opt* file.

To edit your *dsm.opt* file:

1. Group the options into stanzas for each server that your client node contacts for backup, archive, and space management services.
2. Enter a value for each option and remove the leading asterisk (*).
3. Set the DSM_CONFIG environment variable to point to your *dsm.opt* file. For instructions to set this variable, see “Set Environment Variables” on page 34.

Table 7. Setting Space Management Options in Your *dsm.opt* File

Option	Description	Default
<i>optionformat</i>	Specifies the format to use (either standard or short) when you specify HSM client commands. If you change the value, the new value is effective immediately.	Standard
<i>restoremigstate</i>	Restores a file to stubbed (migrated) state if it is backed up after migration or premigration. Tivoli Storage Manager records the migration state of files during backup, so only those files that were migrated or premigrated at the time of backup can be restored to stubbed (migrated) state.	Yes

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

Performing Optional Setup Tasks

In addition to required tasks, an authorized user can perform these optional tasks:

- Assign management classes to your files
- Create an include-exclude options file
- Set environment variables.

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

Assign Management Classes to your Files

Your Tivoli Storage Manager administrator defines management classes that contain specific requirements for migrating files to storage. You assign these management classes to files on your local file systems. The management class that

you assign to a file determines file eligibility for migration. Use the default management class for some or all of your files. Assign different management classes to specific files or groups of files using one or more include statements in your include-exclude options file.

A management class can contain a backup copy group and an archive copy group. Copy groups contain attributes that control the generation, destination, and expiration of backup versions of files and archived copies of files. For information about backup and archive copy groups, see *IBM Tivoli Storage Manager for UNIX: Backup-Archive Clients Installation and User's Guide*. Table 8 on page 31 lists the space management attributes and their defaults that might be included in a management class.

Table 8. Space Management Attributes in a Management Class

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

Table 9 defines some tasks you might perform and the type of management class to assign to your files.

Table 9. Assigning Management Classes to Your Files

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

Display Information About Management Classes

You can display information about management classes for the default migration server only. If your client node contacts one or more additional servers for backup and archive services, use the backup-archive client graphical user interface or the **dsmc query mgmtclass** command to display information about available management classes on those servers.

To display information about management classes that you can assign to your files, follow these steps:

1. Click **Utilities**→**Display Policy Information**.
2. Click the management class for which you want to display information.
3. Click **Help** to display descriptions of management class attributes.
4. Click **Close** to close the Display Policy Information window.

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

```
dsmmigquery -mgmtclass -detail
```

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

Modify your Include-Exclude Options File

You might want to keep certain files, such as system files or files that the HSM client creates and uses, on your local file system at all times. Or, you might want to

include certain files for backup or migration. Set the statements in your include-exclude options file to exclude or include specific files from space management and assign specific management classes to these files. If you do not create an include-exclude options file, all files are considered for backup services and the default management class is used.

For AIX GPFS and Linux86 GPFS file systems only: Do not use the *inclexcl* option with the AIX GPFS or Linux86 GPFS failover environment. Unlike *dsm.opt* and *dsm.sys* files, include-exclude options files are not shared between different nodes of a failover group. Add the include-exclude list directly to the *dsm.sys* file, or verify that the include-exclude list files match on all nodes participating in the local failover group or node set.

For all other file systems: Use the *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 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 use the options in Table 10 to include or exclude files from either backup or migration:

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

Table 10. Include and Exclude Statements

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

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

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

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

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

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

Create an Include-Exclude List

To create an include-exclude list from the graphical user interface:

1. Determine your include and exclude requirements.
2. Select the **Utilities** menu → **Setup Wizard**.
3. Check the **Help me configure the TSM Backup Archive Client** box.
4. Select **Update my options file**.
5. Click **Next** several times until the Domain and include/exclude lists screen displays.
6. Edit your include-exclude list.

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

1. Determine your include and exclude requirements.
2. Create a file. Store your file in any directory to which all users on your workstation have read access.
3. Enter your include and exclude statements.
4. Enter the *inclexcl* option and the name of your include-exclude options file in your dsm.sys file.

Set Environment Variables

Note: You cannot specify the root directory for DSM_DIR, DSM_CONFIG, or DSM_LOG.

Set these environment variables to point to files that Tivoli Storage Manager uses. Use the LANG environment variable to specify the language that you want to use.

Table 11. Environment Variables

Variable	Description
DSM_DIR	Points to the resource files, the dsm.sys file, and the executable file, dsmtca. If you do not set DSM_DIR, the HSM client searches for the executable files in the installation directory.
DSM_CONFIG	Points to your dsm.opt file. <ul style="list-style-type: none"> • If you do not set DSM_CONFIG, the HSM client searches for the options file in the directory to which DSM_DIR points. • If you do not set DSM_DIR, the HSM client searches for the options file in the installation directory. • For Solaris operating systems, dsm.sys and dsm.opt are symbolic links to the actual files that you store in /usr/bin. This protects your files from deletion if you uninstall Tivoli Storage Manager.

Table 11. Environment Variables (continued)

Variable	Description
DSM_LOG	<p>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.</p> <ul style="list-style-type: none">• If you define DSM_DIR but you do not define DSM_LOG, messages are written to dsmerror.log in the directory that you specified in DSM_DIR.• If you do not define DSM_LOG or DSM_DIR, error messages are written to dsmerror.log in the current directory. You receive a warning message if Tivoli Storage Manager cannot write messages to the log file. Processing continues.

Set Bourne and Korn Shell Variables

To set the Bourne or Korn shell, enter the environment variables in the .profile file in your \$HOME directory. For example:

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

The /home/davehil/dsm.opt path identifies the path and file name for your dsm.opt file. The /home/davehil path identifies the directory where you want to store the dsmerror.log file, the executable file, the resource files, and the dsm.sys file.

Set 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/davehil/dsm.opt  
setenv DSM_LOG /home/davehil
```

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

Set Up LAN-free Data Transfer for HSM on AIX GPFS

The HSM on AIX GPFS 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 you to migration and recall 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

To enable LAN-Free support:

- A Tivoli Storage Manager version 5.2.2 or higher client and server is required.
- You must install and configure the Tivoli Storage Manager Managed System for SAN Storage Agent on the client workstation. For more information, refer to the IBM Tivoli Storage Manager for AIX Storage Agent User's Guide, GC32-0771.

LAN-free options

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

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*.

lanfreetcport

Specifies the TCP/IP port number where the Storage Agent is listening. Set this option to *1530*.

See the *IBM Tivoli Storage Manager for UNIX Backup-Archive Clients Installation and User's Guide* for more information about these options.

Use the ***include*** and ***exclude*** options to control LAN-free data transfer, as follows:

Task Assuming that */hsm1* is an HSM managed file system, include the files in the */hsm1/clientdata/lanfree* directory for LAN-free data transfer and assign these files to a LAN-free enabled management class.

Include statement: `include.lanfree /hsm1/clientdata/lanfree
lanfreemgmtclass`

Task Assuming that */hsm1* is an HSM managed file system and the default management class is LAN-free, exclude the files in the */hsm1/clientdata* directory from LAN-free data transfer.

Exclude statement: `exclude /hsm1/clientdata/*`

Set Up Space Management for AIX in an HACMP Environment

Attention: The HACMP environment is valid on AIX JFS file systems only.

The HSM client permits your HSM-managed file system to become a part of a resource group in an HACMP cluster so that if there is a system failure, you can obtain access to your data from another system. To set up space management in an HACMP environment, each node in a cluster, willing to take over HSM-managed file systems, must specify a server stanza for every Tivoli Storage Manager server that your file systems use in a cluster that it might want to manage. Place this stanza in your *dsm.sys* file. You can use a server stanza for more than one file system.

Notes:

1. The HACMP setup in conjunction with HSM is only supported with AIX 5.1.
2. For more information about the HACMP environment for AIX, see *IBM Tivoli Storage Manager for UNIX: Backup-Archive Clients Installation and User's Guide*.

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

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

- The server stanza must point to the same server on each system.

The HSM for HACMP also supports cascading and rotating takeover relationships in the same manner as the backup-archive client. The password handling is the same as that of the HACMP backup-archive client. The `dsm.opt` file and the `dsm.sys` file can reside in the standard `/usr/tivoli` folder, or you can use the `DSM_CONFIG` environment variable that you set in the start script. However, the behavior must be identical on all nodes.

Limitations of HSM Support for the HACMP Environment

Included here are the limitations of HSM support for the HACMP environment.

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

The HSM support for the HACMP environment is not completely integrated with the backup-archive client support for the HACMP environment. For example, the HSM client refers to the `/etc/adsm/SpaceMan/config/dsmmigfstab` file to determine which server to use for a file system. The server might contact a different server for each file system. In contrast, the backup-archive client determines which server to use from the `dsm.opt` file, the `dsm.sys` file, or from a parameter that you specify on the command line when you start a backup-archive command line client. A backup-archive client process might back up, archive, restore, or retrieve from one server. If you need backup-archive client services for different servers, start a new backup-archive client process.

The HSM GUI does not support HACMP commands. After you enter the **`dsmmigfs import`** command, run the **`dsmreconcile`** command for your file system to update the status file.

Start and Stop HSM in an HACMP Environment

The following examples of scripts demonstrate how to start and stop HSM in an HACMP environment on an AIX file system. These scripts are only examples. You must create your own scripts for your environment.

```

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

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

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

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

# killing all running demons to reread the option files
`echo "Killing dsmrecauld ..." >> $LOG`;
`kill -15 \$(ps -aef|grep dsmrecauld |grep -v grep |awk '{print \$2}');
2>&1 1>> $LOG`;
`echo "Killing dsmonitord ..." >> $LOG`;
`kill -15 \$(ps -aef|grep dsmonitord |grep -v grep |awk '{print \$2}');
2>&1 1>> $LOG`;
`echo "Killing dsmscoudt ..." >> $LOG`;
`kill -15 \$(ps -aef|grep dsmscoudt |grep -v grep |awk '{print \$2}');
2>&1 1>> $LOG`;

# starting the demons with the right dsm.sys and dsm.opt
`echo "Starting dsmonitord ..." >> $LOG`;
`echo \$DSM_DIR \$DSM_CONFIG >> $LOG;
dsmonitord 2>&1 1>> $LOG`;
`echo "Starting dsmrecauld ..." >> $LOG`;
`echo \$DSM_DIR \$DSM_CONFIG >> $LOG; dsmrecauld 2>&1 1>> $LOG`;
`echo "Starting dsmscoudt ..." >> $LOG`;
`echo \$DSM_DIR \$DSM_CONFIG >> $LOG; dsmscoudt 2>&1 1>> $LOG`;

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

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

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

```

Figure 2. start_HSM : A Sample Script That Starts the HSM and Imports all File Systems

```

#!/usr/bin/perl
#----- #
call with
#   stop_HSM [filesystems]
#-----
my $LOG="/tmp/HSM.log"; # place of your logfile

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

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

# before you can export the filesystems, you have to stop all demons
`echo "Killing dsmreca1ld ..." >> $LOG`;
`kill -15 \$(ps -aef|grep dsmreca1ld |grep -v grep |awk '{print \$2}');
2>&1 1>> $LOG`;
`echo "Killing dsmonitord ..." >> $LOG`;
`kill -15 \$(ps -aef|grep dsmonitord |grep -v grep |awk '{print \$2}');
2>&1 1>> $LOG`;
`echo "Killing dsmscoutd ..." >> $LOG`;
`kill -15 \$(ps -aef|grep dsmscoutd |grep -v grep |awk '{print \$2}');
2>&1 1>> $LOG`;

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

# to export the filesystems, you have to stop all processes using this filesystem
`echo "Killing the rest ..." >> $LOG`;
for(my $i=0; $i < @FS; $i++){
    @PROCS=split (" ",`fuser -c $FS[$i] 2>/dev/null`);
    while(my $PROC=shift @PROCS){
        $PROC=~s/c//g;
        `kill -9 $PROC 2>&1 1>> $LOG`;
    }
}

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

```

Figure 3. stop_HSM : Stops HSM and Exports All File Systems

Set Up Space Management for Linux86 in an IBM Cluster 1350 Environment

The HSM client permits your HSM-managed file system to become a part of a resource group in a resource group in the cluster so that if there is a system failure, you can obtain access to your data from another system. To set up space management in an IBM Cluster 1350 environment, each node in a cluster, willing to take over HSM-managed file systems, must specify a server stanza for every Tivoli Storage Manager server that your file systems use in a cluster that it might want to manage. Place this stanza in your dsm.sys file. You can use a server stanza for more than one file system.

For more information about the IBM Cluster 1350 environment for Linux, see IBM Redbook *Linux Clustering with CMS and GPFS*, SG24-6601 for more information:
<http://www.redbooks.ibm.com/redbooks/pdfs/sg246601.pdf>

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

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

The HSM for IBM Cluster 1350 also supports cascading and rotating takeover relationships in the same manner as the backup-archive client. The `dsm.opt` file and the `dsm.sys` file can reside in the standard `/opt/tivoli` folder, or you can use the `DSM_CONFIG` environment variable that you set in the start script. However, the behavior must be identical on all nodes.

Limitations of HSM Support for the IBM Cluster 1350 Environment

Included here are the limitations of HSM support for the IBM Cluster 1350 environment.

- The management class information is for the default migration server only.
- The server options information is for the default migration server only.
- Every IBM Cluster 1350 machine must run the same HSM version.

The HSM support for the IBM Cluster 1350 environment is not completely integrated with the backup-archive client support for the IBM Cluster 1350 environment. For example, the HSM client refers to the `/etc/adsm/SpaceMan/config/dsmmigfstab` file to determine which server to use for a file system. The server might contact a different server for each file system. In contrast, the backup-archive client determines which server to use from the `dsm.opt` file, the `dsm.sys` file, or from a parameter that you specify on the command line when you start a backup-archive command line client. A backup-archive client process might back up, archive, restore, or retrieve from one server. If you need backup-archive client services for different servers, start a new backup-archive client process.

The HSM GUI does not support IBM Cluster 1350 commands. After you enter the `dsmmigfs import` command, run the `dsmreconcile` command for your file system to update the status file.

Start and Stop HSM in an IBM Cluster 1350 Environment

The following examples of scripts demonstrate how to start and stop HSM in IBM Cluster 1350 environment on a Linux86 GPFS file system. These scripts are only examples. You must create your own scripts for your environment.

```

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

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

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

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

# killing all running demons to reread the option files
`echo "Killing dsmrecauld ..." >> $LOG`;
`kill -15 \$(ps -aef|grep dsmrecauld |grep -v grep |awk '{print \$2}');
2>&1 1>> $LOG`;
`echo "Killing dsmonitor" >> $LOG`;
`kill -15 \$(ps -aef|grep dsmonitor |grep -v grep |awk '{print \$2}');
2>&1 1>> $LOG`;
`echo "Killing dsmscoutd ..." >> $LOG`;
`kill -15 \$(ps -aef|grep dsmscoutd |grep -v grep |awk '{print \$2}');
2>&1 1>> $LOG`;

# starting the demons with the right dsm.sys and dsm.opt
`echo "Starting dsmonitor" >> $LOG`;
`echo \$DSM_DIR \$DSM_CONFIG >> $LOG`;
dsmonitor 2>&1 1>> $LOG`;
`echo "Starting dsmrecauld ..." >> $LOG`;
`echo \$DSM_DIR \$DSM_CONFIG >> $LOG; dsmrecauld 2>&1 1>> $LOG`;
`echo "Starting dsmscoutd ..." >> $LOG`;
`echo \$DSM_DIR \$DSM_CONFIG >> $LOG; dsmscoutd 2>&1 1>> $LOG`;

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

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

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

```

Figure 4. start_HSM : A Sample Script That Starts the HSM and Imports all File Systems

```

#!/usr/bin/perl
#----- #
call with
#   stop_HSM [filesystems]
#-----
my $LOG="/tmp/HSM.log"; # place of your logfile

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

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

# before you can export the filesystems, you have to stop all demons
`echo "Killing dsmrecauld ..." >> $LOG`;
`kill -15 \$(ps -aef|grep dsmrecauld |grep -v grep |awk '{print \$2}');
2>&1 1>> $LOG`;
`echo "Killing dsmonitord ..." >> $LOG`;
`kill -15 \$(ps -aef|grep dsmonitord |grep -v grep |awk '{print \$2}');
2>&1 1>> $LOG`;
`echo "Killing dsmscouth ..." >> $LOG`;
`kill -15 \$(ps -aef|grep dsmscouth |grep -v grep |awk '{print \$2}');
2>&1 1>> $LOG`;

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

# to export the filesystems, you have to stop all processes using this filesystem
`echo "Killing the rest ..." >> $LOG`;
for(my $i=0; $i < @FS; $i++){
    @PROCS=split (" ",`fuser -c $FS[$i] 2>/dev/null`);
    while(my $PROC=shift @PROCS){
        $PROC=~s/c//g;
        `kill -9 $PROC 2>&1 1>> $LOG`;
    }
}

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

```

Figure 5. stop_HSM : Stops HSM and Exports All File Systems

Chapter 4. Managing Your File Systems

Note: To perform tasks in this chapter, you must have root user authority.

This chapter describes how to:

- Mount your file systems
- Add space management to your local file systems
- Select and update space management settings
- Deactivate and reactivate space management for your local file systems
- Globally deactivate and reactivate space management for your client node
- Remove space management from your local file systems.

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

- Move stub files as native stub files (without the migrated data) to other file systems, whether or not they are space-managed.
- Use the Tivoli Storage Manager raw device backup utility or any other image-based backup utility to restore a space-managed file system to another system other than the original file system.
- Use space-managed file systems within cluster replication tools.
- Use any other tools that transfer images between systems.

Mounting Your File Systems First

Note: For AIX GPFS, Linux86 GPFS, Solaris VxFS, and HP-UX VxFS file systems, an additional mount is not necessary.

Before you can migrate files to Tivoli Storage Manager storage, first mount your file systems either automatically or manually, and then add space management.

For AIX JFS file systems: Mount your file systems twice. Enter these commands to mount them:

```
mount /fs
mount -v fsm /fs
```

The first mount is for the physical AIX JFS file system. The second mount (with type fsm) is for the space management file system.

Mount the file system migrator (FSM) *over* the AIX JFS file system *before* you add space management services. The file system migrator is a kernel extension that intercepts all file system operations and provides any space management support that you require. If you do not mount the FSM, anyone with permission to access migrated files can access and modify your stub files. If a stub file is modified, the HSM client cannot access the migrated file.

After you add space management to an AIX JFS file system, do not unmount the FSM and directly access the native file system. If you unmount the FSM from a file system, enter the following command to remount it:

```
mount -v fsm /fs
```

The `/etc/filesystems` file contains a stanza for each file system to which you add space management. The HSM client updates these stanzas so that the file system and the FSM both mount automatically whenever you restart your system.

Adding Space Management to Your File Systems

Add space management to all file systems for which you need space management services. When you add space management to your file system, the HSM client:

- Creates a hidden directory for your file system named `.SpaceMan` that stores certain information objects required to perform space management tasks. See Table 21 on page 161 for more information about the `.SpaceMan` directory.
- Adds an entry to the `dsmmigfstab` file that includes the file system name to which you added space management and selected your space management settings. The `dsmmigfstab` file is stored in the `/etc/adsm/SpaceMan/config` directory. See Appendix B, “The `.SpaceMan` Directory,” on page 161 for an example of a `dsmmigfstab` file.

Note: Maintain a current backup version of the `dsmmigfstab` file at all times.

- Updates mount information for the native file system.

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

For AIX JFS file systems: If the `/etc/filesystems` file contains a stanza for your file system, the HSM client performs the following:

1. Changes the `mount` attribute from `true` to `false`. This prevents the operating system from mounting your file system automatically when the operating system starts.
2. Adds an `adsmfsm` attribute to the stanza for your file system with a value of `true`. The `adsmfsm` attribute indicates that the HSM client manages the file system.

- **For AIX JFS file systems:** Mounts the file system migrator (FSM) over the native file system.
- Starts space management for your file system.

For AIX GPFS and Linux86 GPFS file systems: The `dsmwatchd` daemon starts at system startup (`inittab` entry). Ensure that DMAPI is enabled on all GPFS file systems that the HSM client manages. Enter the following command to query this information:

```
For AIX GPFS:  
/usr/lpp/mmfs/bin/mmfsfs DevicePath -z
```

```
For Linux86 GPFS:  
/opt/lpp/mmfs/bin/mmfsfs DevicePath -z
```

If it is required, change the value to:

```
For AIX GPFS:  
/usr/lpp/mmfs/bin/mmchfs DevicePath -z yes
```

```
For Linux86 GPFS:  
/opt/lpp/mmfs/bin/mmchfs DevicePath -z yes
```

Note: When DMAPI is enabled on a GPFS file system, the file system can be mounted *only if* a `dsmrecalld` daemon is already set up on one of the cluster nodes within the GPFS node set.

The AIX or Linux86 cluster node to which you add a GPFS file system becomes the preferred node for your file system. If several HSM clients on several AIX or

Linux86 cluster nodes are candidates for managing a GPFS file system, the preferred node has precedence. In the event of failure and takeover by another node within the GPFS node set, the HSM client that manages the file system is not the HSM client on the preferred node. A GPFS node set is a set of AIX or Linux86 cluster nodes that can mount a defined group of GPFS file systems.

Attention: During the mount process and while the HSM client is adding space management to your file systems, do not attempt to access any files in your file systems or perform any tasks against your file systems.

Select Space Management Settings

When you add space management to your file systems, select space management settings that will control the following:

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

This information is stored in the `dsmmigfstab` file located in `config/dsmmigfstab`. You can modify the settings in the `dsmmigfstab` file using the **`dsmmigfs`** command (see “`dsmmigfs`” on page 124). You can also use the **`dsmmigfs`** command to display the settings in this file (see “`dsmmigfs`” on page 131).

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.

The following sections provide information to help you select space management settings for your file systems.

Minimum File Size for Migration

The HSM client does not migrate a file unless doing so saves space on your local file system. For AIX JFS file systems, before a file is considered for migration, the size must be greater than both the stub file size plus one byte and the file system fragment size. If the **`minmigfilesize`** option is defined, the file must be greater than this option value. For example, if the stub file size is 2047 bytes, and the defined block size for a file system is 4096 bytes, the file size must be greater than 4096 bytes before it is eligible for migration.

For AIX GPFS, Linux86 GPFS, Solaris VxFS, and HP-UX file systems, the file size must be greater than both the stub file size and the file system block size before it is eligible for migration.

Set the High and Low Threshold Percentages

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

The percentage that you specify for a low threshold must be the same as, or lower than, the percentage that you specify for a high threshold. For example, to stop migrating files when there are 20 MB of available free space on your file system, set the low threshold to 80 percent.

Note: If the high threshold and the low threshold are the same, space usage must exceed the low threshold before threshold migration begins.

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

Set the 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 percentage that you specify for the low threshold and the percentage that you specify for the high threshold. The default percentage premigrates enough files to make the next occurrence of threshold migration faster. For example, if the high threshold is 90 percent and the low threshold is 80 percent, the premigration percentage is 10 percent. When space usage drops to 80 percent, additional files premigrate until at least 10 percent of the occupied space on your file system contains premigrated files that are listed at the beginning of the current migration candidates list. The next time threshold migration is required, the HSM client replaces those files with stub files on your local file system. This quickly reduces space usage to a low threshold without requiring additional time to copy the files to Tivoli Storage Manager storage.

If demand migration is required, having your eligible files already premigrated hastens that process as well. The HSM client quickly releases at least ten percent of

the space on your local file system, and migrates any additional files that are necessary to return space usage to the low threshold.

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

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 zero. The HSM client does not premigrate any files after threshold or demand migration completes.

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

Set the Age and Size Factors

Set the age and size factors to determine the order in which eligible files migrate to Tivoli Storage Manager storage during automatic migration. The age factor determines how much consideration is given to the number of days since a file was last accessed. The size factor determines how much consideration is given to the size of the file. Files are prioritized for automatic migration in descending order by their priority score. Set a value from zero through 999999999 for both age and size factors. The default is one for both factors.

Table 12 displays how to prioritize your files for migration. If a resident file and a premigrated file have the same priority score, the HSM client assigns the premigrated file a higher priority.

Table 12. Setting the Age and Size Factors

Prioritize Your Files	Set the Factors
Prioritize your files for migration by age only.	Set the age factor to one and the size factor to zero.
Prioritize your files for migration by size only.	Set the age factor to zero and the size factor to one.
Prioritize your files for migration according to their age.	Increase the age factor.
Prioritize your files for migration according to their size.	Increase the size factor.
Prioritize your files for migration in order by age, and within age, by size.	Set the size factor to one, and set the age factor to the size of the largest file that you want to store, or to the size (in 1 KB blocks) of the file system itself. The HSM client assigns the highest priority to the oldest files. If two or more files are the same age, files are prioritized by size.
Prioritize your files for migration in order by size, and within size, by age.	Set the age factor to one, and set the size factor to the number of days equal to the longest life span of a file on your file system. The HSM client assigns the highest priority to the largest files. If two or more files are the same size, files are prioritized by age.

Set the Quota

The quota that you set determines the maximum number of megabytes of data that you can migrate and premigrate from your file system to storage. Specify a quota value from zero through 999999999. The default is the number of megabytes that are assigned for your file system. For example, if 20 gigabytes are assigned for

your file system, the HSM client migrates and premigrates your files from that file system until the total number of megabytes that migrate and premigrate equals 20 gigabytes.

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.

If you set the quota to zero 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. 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.

Set the Maximum Number of Candidates

Because dsmscoutd does not perform reconciliations, this is the maximum number of migration candidates that the dsmscout daemon searches for during a scan of your file systems. When this number is reached, the dsmscout daemon stops. The range of values that you can enter are zero through 9999999. The default is 10000.

Set the Stub File Size

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

For HSM clients on AIX JFS file systems, the minimum stub file size is 511 bytes and the default size is 4095. You can select a stub file size that is smaller than your file system block size. However, this is an inefficient use of space on your file system. Although the HSM client uses only the number of bytes that you specify for stub files, your operating system allocates one block for each stub file. The excess number of bytes remain unused.

For HSM clients on Solaris VxFS and HP-UX VxFS file systems, valid stub file sizes are a multiple of the physical blocksize. For HSM clients on AIX GPFS and Linux86 GPFS file systems, valid stub file sizes are zero, or a multiple of the physical blocksize. The minimum physical block size is 512 bytes. The default logical block size is 4096 bytes.

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

Add Space Management

To add space management to one or more file systems, follow these steps:

1. Click **Space Manager** in the Hierarchical Storage Management window.
2. Click the file systems to which you want to add space management. Do not add space management to the following file systems:
 - root (/)
 - /tmp
 - /usr
 - /var

These file systems contain files that your operating system frequently uses.

3. Click **Selected**→**Add Space Management**. If you select more than one file system, a window for each file system displays. Select the appropriate window.
4. Select values for each space management setting. For information, highlight the setting and press **F1** for help.
5. Click **Add** to add space management to your local file system.
6. Click **OK** when the Information Dialog window confirms that space management is added to your file system.

Add Space Management from the Command Line

To add space management to your local file systems from the command line, use the **dsmmigfs** command. For example, to add space management to the /home file system and set a high threshold of 85 percent and a low threshold of 75 percent, enter the following:

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

See “dsmmigfs” on page 124 for more information about this command.

Add Space Management to Nested File Systems

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

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

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

To add space management to a parent file system and any nested file systems, follow these steps:

1. Unmount the nested file systems.
2. Add space management to the parent file system.
3. Remount the nested file systems.
4. Add space management to each nested file system. See “Adding Space Management to Your File Systems” on page 44 for instructions to add space management to your local file system.

To add space management to a *parent* file system but not to its nested file system, perform steps one through three above.

Mount the Parent File System First For AIX JFS file systems only

For an AIX JFS-nested file system that mounts automatically when you restart your system, follow these steps to mount the parent file system before you mount the nested file system:

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

```
mount /test/migfs1
```

Where `/test/migfs1` is the name of the nested file system.

Add Space Management to an Exported File System

Note: With the exception of the Solaris HSM, the NFS is the only network file system interface to export a file system that HSM supports.

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, enter:

```
/usr/sbin/showmount -a
```
3. Add space management to your file system using the **`dsmmigfs add`** command.
4. To export the NFS file system again, enter:

```
/usr/etc/exportfs -a
```
5. Instruct all NFS clients to mount the exported NFS file system again.

Managing Your File Systems

After you add space management, you can manage your file systems easily. This section describes how to:

- Update space management settings for your file systems.
- Deactivate and reactivate space management on your local file systems.
- Globally deactivate and reactivate space management on your client node.
- Remove space management from your local file systems.

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

- If you change the high and low thresholds or the premigration percentages for a file system, the new values take effect immediately.
- If you change the stub file sizes, the new values take effect only for files that migrate after you make the changes.
- If you change the age and size factors, the new values are used the next time you build a migration candidates list.
- If you reduce the quota, and the data that you currently migrate and premigrate exceeds the new quota, any additional files from your file system do not migrate. Sufficient files must be recalled during automatic or selective recall to drop the total number of megabytes of migrated and premigrated data below the new quota.

Update Space Management Settings

Update space management settings for your file system at any time. Space management can be active, inactive, or globally inactive for a file system. To update space management settings for your file system, follow these steps:

1. Click **Space Manager** in the Hierarchical Storage Management window.
2. Select the file systems that require updating.
3. Click **Selected**→**Update Space Management**. The Update Space Management window displays. When you select more than one file system, a window for each file system displays and arranges them in a cascade. Select the required window to view its contents.
4. Select new values for one or more space management settings. For information about a setting, highlight the setting, and press **F1** for help.

5. Click **Update** to update your space management settings.
6. Click **OK** when the Information Dialog window confirms that your space management settings are updated.

Update Your Settings from the Command Line

To update space management settings for your file system from the command line, use the **dsmmigfs** command. For example, to change the age factor to five and the size factor to one for the /home file system, enter:

```
dsmmigfs update -a=5 -si=1 /home
```

See “dsmmigfs” on page 124 for more information about this command.

Deactivate Space Management from Your File Systems

Deactivate space management from your file systems to temporarily prevent migration, recall, or reconciliation processes from occurring. Any migration, recall, or reconciliation process that currently is in progress completes first. When you reactivate space management on your file system, all space management services resume.

Deactivating space management from your file system:

- Updates space management settings on your file systems
- Adds space management to additional file systems
- Displays information about your file systems while you continue to access resident and premigrated files.

To deactivate space management from your file systems, follow these steps:

1. Click **Space Manager** in the Hierarchical Storage Management window.
2. Select the file systems from space management that you want to deactivate.
3. Click **Selected**→**Deactivate Space Management**.
4. Click **OK** when the Information Dialog window confirms that space management is deactivated from your selected file systems.

Deactivate Space Management from the Command Line

To deactivate space management from your file systems from the command line, use the **dsmmigfs** command. For example:

```
dsmmigfs deactivate /home
```

See “dsmmigfs” on page 127 for more information about this command.

Globally Deactivate Space Management

You can temporarily deactivate space management from all file systems on your client node. When you globally deactivate space management, you cannot perform migration, recall, or reconciliation on any file system. However, any migration, recall, or reconciliation process that is in progress completes before space management is globally deactivated.

Globally deactivate space management before you perform a system maintenance task or before your Tivoli Storage Manager administrator plans to export migrated files from one server and import them to another. When your Tivoli Storage Manager administrator completes the export and import processes, globally reactivate space management on the file systems that you previously deactivated.

Note: If your administrator exports migrated files from one server and imports them to another, update the dsm.sys file so that the client node contacts the

new server for space management services. The administrator can use a **lock node** command to prevent the client node from migrating or recalling files prior to performing the import and export tasks.

Globally deactivating space management:

- Updates space management settings for your file systems
- Adds space management to additional file systems
- Displays information about your file systems while you continue to access resident and premigrated files.

To globally deactivate space management, follow these steps:

1. Click **Space Manager** in the Hierarchical Storage Management window.
2. Click **File**→**Global Deactivate**. An Information Dialog window displays to confirm that the HSM client deactivated space management on your client node.
3. Click **OK** to close the window.

Globally Deactivate HSM from the Command Line

To globally deactivate space management on your client node from the command line, enter:

```
dsmmigfs globaldeactivate
```

See “dsmmigfs” on page 133 for more information about this command.

Reactivate Space Management on Your File Systems

Once you deactivate space management from your file system, you can reactivate it at any time. To reactivate space management, follow these steps:

1. Click **Space Manager** in the Hierarchical Storage Management window.
2. Highlight the file systems for which you want to reactivate space management.
3. Click **Selected**→**Reactivate Space Management**.
4. Click **OK** when the Information Dialog window confirms that space management is reactivated on the selected file systems.

Reactivate Space Management from the Command Line

Use the **dsmmigfs** command to reactivate space management for your file systems from the command line. For example:

```
dsmmigfs reactivate /home
```

See “dsmmigfs” on page 127 for more information about this command.

Globally Reactivate Space Management

You can globally reactivate space management for your client node at any time. All file systems to which you added space management return to their previous states. For example, if you activated space management for a file system when you globally deactivated space management for your client node, it becomes active again. Or, if you deactivated space management for a file system when you globally deactivated space management for your client node, it remains inactive.

To globally reactivate space management for your client node, follow these steps:

1. Click **Space Manager** in the Hierarchical Storage Management window.
2. Click **File**→**Global Reactivate**. An Information Dialog window displays to confirm that space management for the client node is reactivated.
3. Click **OK** to close the window.

Globally Reactivate HSM from the Command Line

Use the **dsmmigfs** command to globally reactivate space management for your client node. For example, enter:

```
dsmmigfs globalreactivate
```

See “dsmmigfs” on page 133 for more information about this command.

Remove Space Management

Space management must be active on your file system to completely remove it. Ensure that you have space on your file system to recall all migrated files. Before you remove space management from your file system, ensure that all activity on your file system has stopped and that the file system is not being accessed.

On an AIX JFS workstation, if the file system is being accessed when you remove space management, the FSM cannot unmount from the file system, and the Tivoli Storage Manager backup-archive and space management functions will not work correctly. If the FSM does not unmount when you remove space management from your file system, enter the **unmount** command for the file system to unmount it after all processes no longer access your file system.

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

- Runs reconciliation for your file system. If any orphaned stub files are located, a notification is posted. Check the orphan.stubs file in the .SpaceMan directory for a list of orphaned stub files. To remove space management, first resolve all orphaned stub files, and then try again. For more information, see “Resolving Orphaned Stub Files” on page 81.
- Determines how much space is required to recall all migrated files. If there is not enough space, the HSM client notifies you. To remove space management, make space available and try again.
- Recalls migrated files to the file system in the most efficient recall order based on where they are stored.
- Notifies the server to delete all migrated files from storage.
- For AIX JFS file systems, unmounts the file system migrator (FSM) from the file system.
- Removes the entry for the file system from the dsmmigfstab file that is stored in the /etc/adsm/SpaceMan/config directory, and deletes the .SpaceMan directory from your file system.
- Updates information for your native file system.

For AIX JFS file systems, if the /etc/filesystems file contains a stanza for the file system, HSM sets the *mount* attribute to true and removes the *adsmfsm* attribute.

For Solaris VxFS file systems, if the /etc/vfstab file contains an entry for the file system, the HSM client updates only the dsmmigfstab file.

To remove space management from one or more file systems, follow these steps:

1. Click **Space Manager** in the Hierarchical Storage Management window.
2. Highlight the file systems that you want to remove from space management.
3. Click **Selected**→**Remove Space Management** to view the Question Dialog window. When you select more than one file system, a Question Dialog window displays for each file system.

4. Click **OK** to recall all migrated files and continue with the remove process. The Remove Status window displays.
5. Click **Cancel** in the Question Dialog window if you want to prevent HSM from removing space management from your file system.
6. Click **Stop** in the Remove Status window if you want to end the process. However, any file that was recalled remains in a premigrated state. Space management is not removed from your file system. Your file system remains in an active state.
7. Click **Close** to close the Remove Status window.

Remove Space Management With a Command

Use the **dsmmigfs** command to remove space management. For example, to remove space management from the /home file system, enter:

```
dsmmigfs remove /home
```

See “dsmmigfs” on page 127 for more information about this command.

Chapter 5. Migrating Your Files

The HSM client automatically migrates files to Tivoli Storage Manager storage when space is needed on your local file system, and automatically recalls files when you access them. Stub files, containing necessary information to locate and recall your migrated files, replace the original files on your local file system.

At any time, you can migrate and recall specific files. For example, if you run threshold migration and some files do not migrate to storage because their last access dates make them ineligible for migration, use selective migration. Files migrate immediately to Tivoli Storage Manager storage. Use selective recall to return them to your local file system.

Note: Any application that touches a file may implicitly cause that file's last access date to change to the time that the application touches it. This is a function of the file system, not the application. Because of this, when the client backs up or archives a file, it may trigger an update to the file's last access date. This can cause problems for other applications such as , whose processing relies on accurate last access dates.

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

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

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

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

For AIX GPFS and Linux86 GPFS file systems only: Newly-created files that you migrate either automatically or selectively must be older than two minutes (two minutes is default) before you can migrate them. Migrating newly-created files of less than five minutes might display incorrect results (resident size) when you use the **dsmdf** and **dsmdu** commands because the GPFS is not synchronized on all nodes when you migrate files. The **dsmdf** command will display correct results after GPFS synchronization, and after the next reconciliation of the file system.

This chapter describes automatic migration, premigration, and selective migration.

Migrating Your Files Automatically

To ensure that free space is available on your local file systems, the HSM client monitors space usage and automatically migrates files whenever it is necessary. Files are prioritized for automatic migration based on the number of days since they were last accessed, their size, and the age and size factors that you set for your file system. The `dsmscoutd` daemon searches these files in cycles and upon request from automatic migration.

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

Demand migration responds to an out-of-space condition on your local file system. The space monitor daemon checks for an out-of-space condition every ten seconds. Threshold migration starts automatically if your file system runs out of space. As files migrate, space becomes available on your file system, and the process that caused the out-of-space condition continues without waiting for threshold migration to complete. You do not receive an out-of-space error message. For example, if you attempt to copy a very large file into your file system, and there is not enough available space for the file, eligible files begin migrating automatically from your local file system to storage. As space becomes available, the process continues to copy the file to your file system.

Premigrate Your Files

For faster migration, the HSM client prepares files for automatic migration using a process called *premigration*. Files are copied to storage beyond the low threshold that you set while the original files remain intact on your local file system. When free space is again needed on your local file system, premigrated files become migrated files. The HSM client verifies that files did not change since they became premigrated. When your premigrated files migrate to storage, stub files replace them on your local file system.

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

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

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

Start Threshold Migration Manually

The space monitor daemon monitors space usage on all file systems for which space management is active. Threshold migration starts automatically to reduce space usage to the low threshold that you set if space usage equals or exceeds the high threshold that you set.

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

To start threshold migration manually for one or more file systems, follow these steps:

1. Click **Space Manager** in the Hierarchical Storage Management window.
2. Click the file systems that you want to select for threshold migration. A window displays for each file system.
3. Select the window that you want to view.
4. Click **Selected**→**Start Threshold Migration**. The Threshold Migration Status window displays.
5. Click **Stop** in the Threshold Migration Status window to stop the migration process. Any files that migrated successfully remain in a migrated state.
6. Click **Close** to close the Threshold Migration Status window.

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

Start Migration With the `dsmautomig` Command

To start threshold migration from the command line, use the `dsmautomig` command. For example, to start threshold migration for the `/home` file system, enter:

```
dsmautomig /home
```

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

See “`dsmautomig`” on page 118 for more information about this command.

Build a New Migration Candidates List

A migration candidates list is a prioritized list of files in your file system that are eligible for automatic migration at the time the list is built. The HSM client uses the list to determine which files to migrate and the order in which to migrate them during threshold and demand migration. Immediately prior to migration, each file is checked again to determine whether it is still eligible for migration.

The `dsmscoutd` daemon continuously scans your managed file systems for candidate requests from the automatic migration processes, and for elapsed scan-interval hours of all your HSM-managed file systems. A slave scout daemon is started to scan for new candidates. After a candidate scan is completed for your file system, a new migration candidates list is created and the slave scout daemon

ends. Another slave scout daemon for the same file system is started after the number of hours you specified on the ***candidatesinterval*** option in your dsm.sys file.

The maximum number of slave scout daemons that can scan in parallel on a system is five. This is the default. You can change this number on the ***maxcandprocs*** option in your dsm.sys file. If a request to build new candidate pools is high, the number of parallel scout daemons might exceed the ***maxcandprocs*** value by three.

When the dsmscoutd daemon locates the maximum number of eligible candidates, it stops scanning your file system and stores the list. Use the ***candidatesinterval*** option in your dsm.sys file to specify how many hours must elapse before the dsmscoutd daemon runs again. Use the ***minmigfilesize*** option in your dsm.sys file to specify the minimum file size, in bytes, for a file to be eligible as a candidate for automigration. The dsmscoutd daemon improves the quality of an existing candidates list during each run.

The dsmscoutd daemon manages two candidates pools (APool and BPool) for each filesystem. These pools are located in /etc/adsm/SpaceMan/candidatesPools. The ***dsmautomig*** process uses the APool, while the dsmscoutd daemon periodically refreshes the BPool. Each pool has a maximum size of 10000 KB each. If the ***dsmautomig*** process migrates all candidates from the APool, the BPool then converts to the APool and the dsmscoutd daemon immediately starts to build a new BPool. Specify the maximum number of migration candidates that the dsmscoutd daemon should search for with the ***maxcandidates*** option on the ***dsmmigfs*** command. Candidates that require more space are not included. The dsmscoutd daemon stores fewer candidates if there is insufficient storage in the parent filesystem of /etc/adsm/SpaceMan/candidatesPools.

Attention: If your system root filesystem does not have sufficient capacity to store this information, either increase capacity, or create a dedicated filesystem with sufficient size mounted to /etc/adsm/SpaceMan/candidatesPools. Kill all running dsmscoutd processes. Create and mount the dedicated file system. Restart the scout daemon manually, using the ***dsmscoutd*** command.

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

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

The HSM client prioritizes files for migration based on the number of days since they were last accessed, their size, and the age and size factors that you set for your file systems. The age factor determines how much consideration is given to the number of days since a file was last accessed, and the size factor determines how much consideration is given to the size of a file.

Migrating Selected Files

Attention: On large file systems, selective migration can take a while.

Use selective migration either from menus or from the command line to move specific files from your local file systems to storage. For example, if you know that you will not be using a particular group of files for an extended time, you can

migrate them to storage to free additional space on your local file system. According to the space management options and settings that you select, the HSM client migrates files that are eligible for selective migration.

When you migrate a file selectively, the access time (atime) for the file does not change.

The number of days since you last accessed a file has no effect on whether your file is eligible for selective migration. This is different from automatic migration. An eligible file must meet the following management class requirements:

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

Migrate Selected Files

To migrate selected files from your local file system to storage:

1. Click **Selective Migration** in the Hierarchical Storage Management window. The Selective Migration window displays.
2. Click **Select File Systems** in the Selective Migration window to select a file system.
3. Click **List Selection** in the Selective Migration window to view a list of the selected files.
4. Select the file systems that contain the files you want to migrate in the Select File Systems for Selective Migration window.
5. Click the **Select** button. A directory tree of the file systems that you selected displays in the Selective Migration window.
6. Select the files that you want to migrate using one of the following methods:
 - To select all files in a directory, click **Select All Files in A Directory**, and select one or more directories.
 - To select all files in all subdirectories, click **Select All Files in All Subdirectories**, and select one or more directories.
 - To select individual files, deselect both the **Select All Files in A Directory** and the **Select All Files in All Directories** check buttons. Within the File List View field of the window, select a directory and select the files that you want to migrate.

By default, the HSM client displays resident and premigrated files in the File List View field only. To display migrated files, click **View options** on the **View** menu.

7. Click **Migrate**. The Selective Migration Status window displays. To stop migrating files, click **Stop** at the bottom of the window. If you click **Stop**, any files already migrated remain migrated.

If you set the *tapeprompt* option to *yes* in your *dsm.opt* file, and the destination for a migrated file is a storage pool that consists of removable media, such as tape, you are prompted to wait for the medium to mount, or to skip the file. See “Assign Management Classes to your Files” on page 29 for information about management class attributes.

8. Click **OK** to close the Information Dialog box after you receive confirmation that all files migrated successfully.
9. Click **Return** to return to the Selective Migration window.
10. Click **Cancel** to close the Selective Migration window.

Start Selective Migration With the **dsmmigrate** Command

To migrate one or more files to storage from the command line, use the **dsmmigrate** command. For example, to migrate a file named `proj1rpt` from the `/home/proja` directory, enter:

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

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

To migrate files in any subdirectory below the specified directory that matches the file specification, use the *-recursive* parameter. For example, to migrate all files in a directory named `/migfs2/test/dir1` and in all of its subdirectories, enter:

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

See “**dsmmigrate**” on page 139 for more information about this command.

Chapter 6. Backing Up and Restoring Files

Backing up your files protects against loss or corruption of your data. Use the Tivoli Storage Manager backup-archive client to back up your files regularly, regardless of whether they are resident, migrated, or premigrated. Back up and restore your migrated and premigrated files in the same manner as you back up and restore files that reside on your local file systems.

See *IBM Tivoli Storage Manager for UNIX: Backup-Archive Clients Installation and User's Guide* for information about backing up and restoring files.

Backing Up, Migrating, and Restoring Your 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.

To restore stub files rather than backup versions of your files, use the **restoremigstate** option with the Tivoli Storage Manager backup-archive client **restore** command. This is useful if one or more of your local file systems is damaged or lost. Your migrated and premigrated files remain intact, and you need only restore your stub files. To recreate stub files for any migrated or premigrated files that are not backed up, use the **dsmmigundelete** command.

Note: If you back up and migrate data to tape, use separate tape drives for backed up and migrated data or make the target storage pools at the top of the primary hierarchy reside on disk (recommended). Otherwise, you cannot back up Tivoli Storage Manager-migrated files. A transparent recall can wait a while for the tape drive to become available. To improve performance, migrate your data to disk storage rather than to tape. Performance requirements for migrated data might be stricter than those for backup data.

If you back up files to one server and migrate them to a *different* server, the HSM client cannot verify that current backup versions of your files exist before you migrate them. Use the backup-archive client to restore the actual backup versions only. You cannot use the backup-archive client to restore stub files for your migrated files. Use the **dsmmigundelete** command to recreate stub files for any migrated or premigrated files that are lost.

Back Up Your Files Before You Migrate Them

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

If you back up files to one server and migrate them to a *different* server, the HSM client cannot verify that current backup versions of your files exist before it

migrates them. Any management class that you assign to files must specify that current backup versions are not required prior to migration. Otherwise, you cannot migrate your files.

Back Up Your Files After You Migrate Them

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 migrate them, files are copied from the migration destination to the backup destination. Files are not recalled to your local file system.

Valid for AIX JFS file systems only:

A backup of a migrated file signals a recall if it has been touched within the last two minutes. The file is premigrated.

If you back up files to a *different* server after you migrate them, the backup-archive client accesses the files using:

- Migrate-on-close recall mode if the recall mode that you set for the file is normal or migrate-on-close. The file resides on your local file system only until a backup version of the file is sent to your backup server.
- Read-without-recall mode if the recall mode that you set for the file is read-without-recall.

For more information about recall modes, see Chapter 5, “Migrating Your Files,” on page 55.

Back Up and Restore Your Premigrated Files

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

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

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

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

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

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

Restore Your Migrated Files

If you back up files after they are migrated to the *same* server, you can restore migrated files using one of the following methods:

- Restore a backup version if the stub file is lost or corrupted and you want to restore the entire file.

Set the **restoremigstate** option to *no* in your `dsm.opt` file. The file becomes a resident file. The migrated copy of your file is removed from Tivoli Storage Manager storage when it expires. Specify the expiration with the **migfileexpiration** option in your `dsm.sys` file.

- Restore a stub file if the stub file on your local file system is lost or corrupted, and you want the file to remain migrated.

Set the **restoremigstate** option to *yes* in your `dsm.opt` file.

For AIX GPFS, Linux86 GPFS, Solaris VxFS, and HP-UX VxFS file systems only: Files with ACLs are restored to a resident state, even when you specify **restoremigstate yes**.

- Recreate a stub file if a backup version of your migrated file does not exist. If a backup version of a migrated file does not exist, you cannot use the backup-archive client to restore a stub file for that file. To recreate a stub file for a migrated file, use the **dsmmigundelete** command.

For more information about these options, see “Migfileexpiration” on page 104 and “Restoremigstate” on page 111.

If you back up files to one server and migrate them to a *different* server, use one of the following methods to restore your migrated files:

- Restore a backup version if the stub file is lost or corrupted. Use the **dsmmigundelete** command. If a migrated copy of your file exists in storage, the migrated copy of the file is removed from storage when it expires. The file becomes a resident file. Specify the expiration with the **migfileexpiration** option in your `dsm.sys` file.
- Recreate a stub file if a backup version of your file does not exist, or if a stub file is lost or corrupted and you want the file to remain migrated. When you back up and migrate files to different servers, you cannot use the backup-archive client to restore a stub file for a migrated file. To recreate a stub file for a migrated file, use the **dsmmigundelete** command.

Recreating Stub Files for Your Migrated Files

If you back up and migrate files to the *same* server, and backup versions of your files exist, use the backup-archive client to restore your stub files.

If you back up and migrate files to a *different* server and your stub files are erased or corrupted, use the **dsmmigundelete** command to recreate them.

Note: When you use the **dsmmigundelete** command to recreate stub files for migrated files, stub files are recreated for all eligible migrated files in the file system that you specify. You cannot use this command to recreate stub files for individual files or specific groups of files.

The **dsmmigundelete** command creates stub files for any premigrated files for which an original file does not exist on your local file system.

- Enter the **dsmmigundelete** command *without* the **expiring** option if reconciliation was not run since the files were deleted. The HSM client performs the following action for the file system that you specify:

File	Description
Migrated files	Recreates a stub file for a migrated file if a corresponding stub file does not exist on your local file system, and the migrated file was not marked for expiration.
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.

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

File	Description
Migrated files	Recreates a stub file for a migrated file if a corresponding stub file does not exist on your local file system, whether the migrated file was marked for expiration or not.
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.

When you enter the **dsmmigundelete** command:

- The **dsmmigundelete** command creates a stub file containing the necessary information to recall the corresponding file from storage. It does not contain any leading bytes of data from the file.
- The recall mode that you previously set for a migrated file is not stored in a recreated stub file. The recall mode for the file is set to normal.
- If a directory path does not exist in your local file system for a migrated or premigrated file, the HSM client does not create a stub file for that file.
- A stub file is recreated with the name of the file at the time it was migrated. If you rename a file after it is migrated, the file name is not updated on the server.
- If you have more than one migrated file in storage with the same name that is marked for expiration, a stub file is created again for the file with the most recent modification time (mtime).
- If the name of a migrated or premigrated file in storage is the same as the name of a file that currently resides on your local file system, the HSM client replaces the file on your local file system with a stub file only if the modification time for the migrated or premigrated file is newer than the modification time for the file on your local file system.
- The HSM client can recreate a stub file for a migrated file, or create a stub file for a premigrated file even if:
 - The file was never backed up.
 - The migrated or premigrated file resides on a different server other than the server on which backup copies of the file reside.
 - The file was migrated or premigrated after the last incremental backup.

Attention: The **dsmmigundelete** command does not support hardlinked files. If you attempt to recreate a stub file for a hardlinked file, a stub file is not recreated *unless* all of the files that are hardlinked together are deleted from your local file system. When one file in a set of hardlinked files is migrated, all of the hardlinked files in the set become stub files. When the **dsmmigundelete** command recreates a stub file for a hardlinked file, the stub file has the same name as the file that was originally migrated. Stub files are not recreated for any other files that were previously in the hardlinked set of files.

For more information about using the **dsmmigundelete** command, see “dsmmigundelete” on page 141.

Restoring Your File Systems

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

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

- If the backup versions and migrated files reside on the *same* server, use the **restore** command and set the **restoremigstate** option to *Yes* (the default). The HSM client restores backup versions of resident files, and restores migrated and premigrated files to stub files.
- If the backup versions and migrated files reside on *different* servers, or if there are no backup versions of migrated and premigrated files, use the **dsmmigundelete** command to recreate stub files for migrated and premigrated files.

Note: When you restore an entire file system, perform this task *before* you restore backup versions of your resident files.

Using one of these methods restores your file system to its state as of the last incremental backup. Premigrated files change to a migrated state.

Back Up and Migrate Files to the Same Server

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

1. Follow your operating system instructions to establish the file system again.
2. Mount the file system. For information about mounting file systems, see “Mounting Your File Systems First” on page 43.
3. If the `/etc/adsm/SpaceMan/config/dsmmigfstab` file contains an entry for your file system, remove the entry from the file. Ensure that the file system does not contain the `.SpaceMan` directory. Otherwise, you will not be able to add space management to your file system.
4. Add space management to your file system. For information about adding space management, see “Adding Space Management to Your File Systems” on page 44.
5. Enter the **dsmc restore** command and set the **restoremigstate** option to *yes* (the default). This command is provided with the backup-archive client. For example, to restore the `/home` file system, enter:

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

Enter the **dsmc restore** command with the **restoremigstate** option set to *yes*. The HSM client restores backup versions of resident files and restores stub files for migrated and premigrated files.

Back Up and Migrate Files to a Different Server

Follow these instructions to restore your file system if you back up files to one server and migrate them to another, or if backup versions are not available for migrated and premigrated files.

1. Follow your operating system instructions to establish the file system again.
2. Mount the file system. For information about mounting file systems, see “Mounting Your File Systems First” on page 43.
3. If the `/etc/adsm/SpaceMan/config/dsmmigfstab` file contains an entry for your file system, remove the entry from the file. Ensure that the file system does not contain the `.SpaceMan` directory. Otherwise, you will not be able to add space management to your file system.
4. Add space management to your file system. For information about adding space management, see “Adding Space Management to Your File Systems” on page 44.
5. Enter the **dsmc restore** command with the **dironly** option to restore the directory structure of your file system. This command is provided with the backup-archive client. For example, to restore the directory structure for the `/home` file system, enter:

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

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

6. Enter the **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, enter:

```
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.

For more information about the **dsmmigundelete** command, see “**dsmmigundelete**” on page 141.

7. Enter the **dsmc restore** command with the **replace** option set to **no** to restore backup versions of previously resident files. For example, to restore backup versions of all the remaining files in the `/home` file system, enter:

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

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

Restoring A Disk

To restore a disk in the event of disk loss, restore your:

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

Restoring Your Operating System

If you can run the backup-archive client, you can recover files. If you can run the HSM client, you can recreate stub files for your migrated files and create stub files for your premigrated files. If you lose the file system that contains your Tivoli Storage Manager clients, you must reinstall the clients before you can recover your files. If you lose the file system that contains the operating system and communications software, you must recover them *before* you can connect to your server. To minimize the impact of such losses, create a set of installation media that will restore your system to a state that permits contact with the server. The installation media should contain:

1. A startable operating system that permits you to perform basic functions.
2. A correctly-configured communications program that permits you to establish communications with the server.
3. A backup-archive client and an HSM client with customized options files. Command-line clients are sufficient.

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.

Restore Your File Systems

Use these instructions to restore your file systems if you back up and migrate files to the *same* server. Enter commands to restore your file systems in the exact order that is presented in the following steps. Otherwise, you might not obtain the results that you want.

1. Reinstall the backup-archive client and the HSM client.
2. Follow your operating system instructions to establish your native file systems again.
3. Mount each file system that you want to restore. For information about mounting file systems, see “Mounting Your File Systems First” on page 43.
4. Add space management to each file system that the HSM client previously managed. For information about adding space management, see “Adding Space Management to Your File Systems” on page 44.
5. Enter the backup-archive client **dsmc restore** command with the **restoremigstate** option set to *yes*. You can perform this step for one or more file systems at the same time.

The backup-archive client restores backup versions of resident files and stub files for migrated and premigrated files. For example, to restore the /home file system, enter:

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

Backup Server and Migration Server are Different

To restore file systems to *different* servers for backup and migration, follow these steps:

1. Reinstall the backup-archive client and the HSM client.
2. Follow your operating system instructions to establish your native file systems again.
3. Mount each file system that you want to restore. For information about mounting file systems, see “Mounting Your File Systems First” on page 43.
4. Add space management to each file system that the HSM client previously managed. For information about adding space management, see “Adding Space Management to Your File Systems” on page 44.
5. Enter the **dsmc restore** command with the **dironly** option to restore the directory structure of each file system. This command is provided with the backup-archive client. You can perform this step at the same time for one or more file systems. For example, to restore the directory structure for the /home file system, enter:

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

When you use the **dironly** option with the **dsmc restore** command, only those backed up directories for the file systems that you specify are restored. For each directory, attributes such as access permissions or, on an AIX workstation, an access control list, are restored.

6. Enter the **dsmmigundelete** command to recreate stub files for migrated files and to create stub files for premigrated files for each file system for which you restored the directory structure. If you enter the **dsmmigundelete** command without a file system specification, the default is all file systems for which space management is active. Space management is automatically activated for a file system when you add space management to it. 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, enter:

```
dsmmigundelete /home
```

Note: When you restore a 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.

For more information about the **dsmmigundelete** command, see “dsmmigundelete” on page 141.

7. Enter the **dsmc restore** command with the **replace** option set to *no* to restore backup versions of previously resident files. For example, to restore backup versions of all remaining files in the /home file system, enter:

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

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

Chapter 7. Archiving and Retrieving Files

You can archive your files at any time and retrieve them to your local file systems when you need them. Use the Tivoli Storage Manager backup-archive client to archive and retrieve copies of your migrated or premigrated files in the same manner as you would archive and retrieve copies of files that reside on your local file systems. See *IBM Tivoli Storage Manager for UNIX: Backup-Archive Clients Installation and User's Guide* for more information about archiving and retrieving files.

Archiving Migrated or Premigrated Files

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

If you archive a copy of a premigrated file to the same server to which it was migrated or to a different server, a copy of the file is sent from your local file system to storage.

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

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

For AIX JFS file systems only: You can archive a copy of a migrated file to a server other than the one to which it migrated. If you do, and you set the recall mode for a file to normal or migrate-on-close, the backup-archive client accesses the file using the migrate-on-close recall mode. The file resides on your local file system until a copy of the file reaches the archive destination.

If you set the recall mode for a file to read-without-recall, the backup-archive client uses that mode during an archive process.

For Solaris VxFS, HP-UX VxFS, AIX GPFS, and Linux86 GPFS file systems only: If you archive a copy of a migrated file to a server other than the one to which it migrated, the backup-archive client accesses and recalls the file. It resides on your local file system in a premigrated state until it is migrated again, or until it receives resident status.

Retrieving Archived Copies of Files

To retrieve an archived copy of a migrated or premigrated file to your local file system, set the ***restoremigstate*** option to *no* in your dsm.opt file. The file is restored to your local file system as a resident file. The migrated copy of your file is removed from storage when it expires.

To retrieve a file to your local file system and maintain a copy of the migrated file in storage, set the ***restoremigstate*** option to *yes* in your dsm.opt file.

For AIX GPFS, Linux86 GPFS, Solaris VxFS, and HP-UX VxFS file systems only: Files with ACLs are restored to a resident state, even when you specify *yes* on the ***restoremigstate*** option in your dsm.opt file.

Chapter 8. Recalling Your Migrated Files

To return a migrated file to your workstation, access the file in the same way as you would access a file that resides on your local file system. The HSM recall daemon automatically recalls the migrated file from Tivoli Storage Manager storage. This process is referred to as transparent recall. If you want to return specific migrated files to your local file system, use selective recall. When you selectively recall a file, the HSM client stores it to its originating file system.

Whenever you recall a file to your local file system, the access time (atime) for the file changes to the current time unless you set the timestamp control mode on the **dsmmode** command to preserve the last recorded access time for your file. This is valid for AIX JFS file systems only.

This chapter discusses both transparent recall and selective recall.

Recalling Your Migrated Files Automatically

Attention: A transparent recall process waits for a tape drive to become available. If you migrate and back up your data to tape, use separate tape drives for both migrated data and backed up data. Otherwise, you cannot back up migrated files. To improve performance, migrate your data to disk storage instead of to tape. Performance requirements for migrated data might be more strict than for backed up data.

Transparent recall automatically returns a migrated file to its originating local file system when you access it. When you recall a file, the HSM client leaves a copy of your recalled file in storage. The file becomes premigrated 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.

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.

Recall Your Files in an AIX JFS Environment

Valid for AIX JFS file systems only

In an AIX JFS environment, the HSM client provides the following two additional recall modes that recall a migrated file to its originating file system:

- Migrate-on-close for a 32-bit and a 64-bit client
- Read-without-recall for 32-bit client only

These recall modes are for read access on migrated files only. You cannot set a recall mode for a resident or a premigrated file.

Attention: Do not use the migrate-on-close or read-without-recall mode for a file that you migrate from a file system that an NFS server exports. Because NFS opens and closes a file many times when an NFS client accesses it, performance can be severely hindered.

To set or change recall modes on your migrated files, use the **dsmatrr** command. The recall mode that you set for a migrated file remains associated with that file only as long as the file remains migrated. See “Use the dsmatrr Command” on page 77 for more information about this command.

If you change the recall mode for a migrated file, or for a specific execution of a recall process, you change how the HSM recall daemon recalls a migrated file. See “Use the dsmmode Command” on page 77 for more information about this command.

If you recreate a deleted stub file, the recall mode for the file becomes normal (the default) despite what you previously set.

You can access a file using more than one process at a time. Another process can cause a file that you are accessing in migrate-on-close or read-without-recall mode to remain on your local file system as a resident or premigrated file. For example, if you set the recall mode to normal (the default) for a migrated file, and you access it with a process using the migrate-on-close recall mode and you do not modify the file, you would expect the file to be in a migrated state when you close it. But if someone else accesses the file at the same time with a process using normal recall mode or with a process that modifies the file in some way, the file remains on your local file system as either a resident or a premigrated file, depending on the actions that the other process takes.

Use the Migrate-on-Close Mode **Valid for AIX JFS file systems only**

The migrate-on-close mode temporarily recalls a migrated file to its originating file system. The recalled file remains on your local file system for as long as it remains open. If you do not modify the file before you close it, the HSM client replaces it with a stub file on your local file system. It then becomes a migrated file. The copy that currently resides in Tivoli Storage Manager storage remains valid because you did not modify the file. If more than one process accesses the file at one time, the HSM client returns the file to a migrated state *only if* all the processes do not modify the file. If any one process modifies the file, it remains on your local file system as a resident file.

Use the **dsmatrr** command to set the recall mode to migrate-on-close for migrated files that you normally read but that you do not modify. When you set the recall mode to migrate-on-close, this mode remains associated with your files until you:

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

For more information about changing recall modes on migrated files, see “Use the dsmatrr Command” on page 77.

Use the Read-Without-Recall Mode **Valid for AIX JFS 4.3 file systems only**

The read-without-recall mode reads a migrated file from Tivoli Storage Manager storage without storing it on your local file system. The HSM client reads information sequentially from the migrated file, and caches that information in a memory buffer on your workstation. This occurs either when the processes that access the file do not modify it, or when the file is executable but the process does not execute the file. If the file is a binary executable file, the file is recalled to your

local file system if a process executes it. The file remains on your local file system as a premigrated file or it is returned to a migrated state, depending on which recall mode the process used that stored it on your local file system.

Attention: Do not set the read-without-recall mode on a file that an NFS server exports. This attribute holds recalled file data in memory. It is not compatible with NFS access, which uses asynchronous reads. If you accidentally set this attribute on a file that a remote NFS client accesses, the HSM client recalls the file and automatically sets the recall attribute of the file to normal.

When you use the read-without-recall mode, the following actions can occur:

- If more than one process accesses a file at the same time, the HSM client reads the file from storage without storing it on your local file system. It does this *only if* all the processes do not modify the file, do not use memory mapping, and do not execute the file if it is a binary executable file. If any one process modifies the file, uses memory mapping, or executes the file, the HSM client copies the file to your local file system. If you do not modify the file, it remains on your local file system as a premigrated file, or it is returned to a migrated state, depending on the recall mode that you used to store it on your local file system. If you modified the file, it remains on your local file system as a resident file.
- If a process that accesses a file also writes to, or modifies the file, the HSM client recalls the file to its originating file system as a resident file. If a file is a binary executable file and it is executed, the file is recalled to its originating file system. If you must return a file to a migrated condition after it is executed, set the recall mode to migrate-on-close for a process that executes the file. Depending on the recall mode that the process uses to store it on your local file system, the file remains on your local file system as a premigrated file, or it is returned to a migrated state.

Use the read-without-recall mode for single access, sequential reads of non-executable files. Accessing a read-without-recall file with more than one process at the same time, or seeking backward in a file can significantly affect the performance of each process that accesses your files. Each of these actions can cause the HSM client to break the connection with the Tivoli Storage Manager server and start the read-without-recall process again. Each read process in progress for the file is delayed until the next piece of information it needs is cached in a memory buffer. Only a small portion of a read-without-recall file is cached in a memory buffer to keep to a minimum the memory overhead of read-without-recall files.

Using the read-without-recall mode to access more than one file can cause a conflict for resources on the Tivoli Storage Manager server. For example, if a single process opens more than one file, and those files reside on the same sequential media, the process obtains access to only one file at a time. If a process reads only the first megabyte of data from each file, and it leaves the files open for additional reads later, one or more of the reads can time out if the process has been waiting too long, and the Tivoli Storage Manager server must free resources to perform the next read.

When you set the recall mode for a migrated file to read-without-recall, that mode remains associated with your file until you:

- Change the recall mode (**dsmatrr** command)
- Access the file with a process that writes to, or modifies the file (**dsmatrr** command)
- Open the file with a process that uses memory mapping

- Recall the file selectively (**dsmrecall** command)
- Recreate the stub file for a migrated file (**dsmmigundelete** command).

If a file is a binary executable file and you run it with execute permission, the recall mode is normal.

Recall Your Files in an AIX GPFS Environment

Use the Partial File Recall Mode

Valid for AIX GPFS HSM only

In an AIX GPFS environment, the HSM client provides the partial file recall mode that recalls a portion of a migrated file. This recall mode is for read access on migrated files that were transferred without compression only. Files that were migrated using compression will always be recalled completely. If a file qualifies for partial file recall (if it is larger than MINPartialrecallsize) *and* compression is also turned on, the partial file recall mode prevails, and the file is sent without compression.

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. When HSM intercepts a read request for such a file and the file is migrated, it will calculate which portion of the file to recall based on the offsets contained in the read request. This results in time and disk space savings, since only a portion of the file is recalled, using less local disk space.

Use the **dsmattr** command to set the recall mode to partial file recall for migrated files that you normally read but that 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

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

- Set the **-RECALLmode** option of the **dsmattr** command to *partialrecall (p)*. This specifies that a file is recalled using partial file recall, regardless of its size. See “**dsmattr**” on page 116 for more information.
- 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. See “**dsmmigfs**” on page 124 for more information.

For more information about changing recall modes on migrated files, see “Use the **dsmattr** Command” on page 77.

Streaming Recall Mode

Valid for AIX GPFS HSM

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

Note: Partial file recall mode takes precedence over streaming recall mode. If a file is smaller than its file system’s MINPartialrecallsize (as configured via the

dsmmigfs – **add or update** commands), or MINPartialrecallsizes is set to 0, normal or streaming recall mode takes precedence.

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

- Set the **-RECALLmode** option of the **dsmattr** command to *Streaming (s)*. This specifies that you want to enable an asynchronous recall of migrated files. See “**dsmattr**” on page 116 for more information.
- Set the **-MINStreamfilesize** option of the **dsmmigfs** command to specify a number to enable or disable an asynchronous recall of migrated files. See “**dsmmigfs**” on page 124 for more information.

How HSM Determines Which Recall Mode to Use

Table 13 displays the recall mode that is used depending on the:

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

Table 13. Recall Modes

If the recall mode for a migrated file is:	And the recall mode for a process is (AIX JFS only):	If the process:	This recall mode is used:
Normal	Normal	Does not modify the file	Normal. File becomes premigrated.
Normal	Normal	Modifies the file	Normal. File becomes resident.
Normal	Migrate-on-close	Does not modify the file	Migrate-on-close. File returned to migrated state.
Normal	Migrate-on-close	Modifies the file	Normal. File becomes resident.
Migrate-on-close	Normal	Does not modify the file	Migrate-on-close. File returned to migrated state.
Migrate-on-close	Normal	Modifies the file	Normal. File becomes resident.
Migrate-on-close	Migrate-on-close	Does not modify the file	Migrate-on-close. File returned to migrated state.
Migrate-on-close	Migrate-on-close	Modifies the file	Normal. File becomes resident.
Read-without-recall	Normal	Does not: <ul style="list-style-type: none"> • Modify the file • Use memory mapping • Execute the file and the file is a binary executable file 	Read-without-recall. File remains migrated.
Read-without-recall	Normal	<ul style="list-style-type: none"> • Modifies the file • Uses memory mapping • Executes the file and the file is a binary executable file 	Normal. If modified, file becomes resident. If not modified, file becomes premigrated.

Table 13. Recall Modes (continued)

If the recall mode for a migrated file is:	And the recall mode for a process is (AIX JFS only):	If the process:	This recall mode is used:
Read-without-recall	Migrate-on-close	Does not: <ul style="list-style-type: none"> • Modify the file • Use memory mapping • Execute the file and the file is a binary executable file 	Read-without-recall. File remains migrated.
Read-without-recall	Migrate-on-close	<ul style="list-style-type: none"> • Uses memory mapping • Executes the file and the file is a binary executable file 	Migrate-on-close. File returns to migrated state.
Read-without-recall	Migrate-on-close	Modifies the file	Normal. File becomes resident.
Partial file recall	Normal	Does not modify the file	Partial file recall. File remains migrated. If the entire file is recalled, the file becomes premigrated.
Partial file recall	Normal	Modifies the file	Normal. File becomes resident.
Streaming	Normal	Does not modify the file	Streaming recall mode. The user application can access the file after its data (in megabytes specified with the <code>MINStreamfilesize</code> option of the <code>dsmmigfs</code> command) have been recalled.
Streaming	Normal	Modifies the file	Normal. File becomes resident.

Recalling Your Migrated Files Selectively

Use selective recall if you want to return specific migrated files to your local file system. The access time (`atime`) changes to the current time when you selectively recall a migrated file.

When you selectively recall a group of files, they are recalled in the most efficient, time-saving order based on where they are stored. For example, if some of your files are stored on a disk storage device and some of your files are stored on a tape storage device, the HSM client recalls all of your files that are stored on the disk storage device first. It next recalls all of your files that are stored on the tape storage device.

For AIX JFS file systems only: Selective recall overrides the recall mode (`migrate-on-close` or `read-without-recall`) that you set for a migrated file. You can selectively recall a file regardless of the mode that you set for your file.

Use Selective Recall

To recall selected files, follow these steps:

1. Click **Selective Recall** in the Hierarchical Storage Management window.

2. Click **Select File Systems** in the Selective Recall window to select a file system.
3. Select the file systems that contain the files you want to recall.
4. Click the **Select** button. A directory of the file systems that you selected displays in the Directory Tree View field in the Selective Recall window.
5. To select all files in a directory, click the **Select All Files in A Directory** button and select one or more directories.
6. To select all files in all subdirectories, click the **Select All Files in All Subdirectories** button and select one or more subdirectories.
7. To select individual files, click **List Selection** to view a list of your selected files. Select the files that you want to recall in the File List View portion of the window. The HSM client displays migrated files in the File List View window only.
8. To display premigrated and resident files, click **View→View options**.
9. Click **Recall**. The Selective Recall Status window displays.
10. Click the **Stop** button at the bottom of the window to stop recalling files. Any recalled files remain on your local file system in a premigrated state.
11. Click **OK** to close the Information Dialog box after you receive confirmation that all files are recalled.
12. Click **Return**. The Selective Recall window displays.
13. Click **Cancel** to close the Selective Recall window.

Using the `dsmrecall` Command

To selectively recall files from storage from the command line, use the **`dsmrecall`** command. For example, to recall a file named `proj1rpt` to the `/home/proja` directory, enter:

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

For AIX JFS file systems only: When you selectively recall a file, you override the migrate-on-close or read-without-recall mode that you previously set for a migrated file.

See “`dsmrecall`” on page 148 for more information about this command.

Use the `dsmattr` Command

Valid for AIX JFS and AIX GPFS only

To set or change the recall mode for one or more migrated files, use the **`dsmattr`** command. Select normal, migrate-on-close, read-without-recall, partialrecall, or streaming recall mode. See “`dsmattr`” on page 116 for more information about this command.

Use the `dsmmode` Command

Valid for AIX JFS file systems only

Use the **`dsmmode`** command to set a recall mode for a process to normal or migrate-on-close. If you set the recall mode for a file to migrate-on-close, and any process modifies the file, it remains on your local file system. When you set the recall mode to normal, the file also remains on your local file system if it is accessed by processes that use the migrate-on-close mode, and at the same time, at

least one process accesses the file or modifies it using normal recall mode. See “Use the dsmmode Command” on page 77 for more information about this command.

Chapter 9. Reconciling Your File Systems

Note: To perform tasks in this chapter, you must have root user authority.

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

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

Starting Automatic Reconciliation

The HSM client automatically reconciles each file system for which space management is active. To specify how often reconciliation runs, modify the setting on the **reconcileinterval** option in your `dsm.sys` file. The default is every 24 hours. To specify how many file systems automatically are reconciled at one time, modify the setting on the **maxreconcileproc** option in your `dsm.sys` file. The default is three file systems.

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

For more information about these options, see Chapter 12, “Using Options,” on page 89.

Table 14 describes the tasks that automatic reconciliation performs on files and file systems.

Table 14. Reconciliation Tasks

Object	Reconciliation Tasks
Migrated files	<ul style="list-style-type: none">• Verifies that a stub file exists on your local file system for each migrated file in storage.• Marks a migrated file for expiration if:<ul style="list-style-type: none">– You deleted the stub file from your local file system– You recalled the file and modified it.• Removes a migrated file from storage if it expired.• Updates the status file.
Premigrated files	<ul style="list-style-type: none">• Verifies that premigrated files are still valid.• Marks a premigrated file for expiration if:<ul style="list-style-type: none">– You deleted the file from your local file system– You accessed the file and modified it.• Removes a copy of a premigrated file from storage if it expired.
Stub files	Records, in the <code>orphan.stubs</code> file, the name of any file for which a stub file exists on your local file system, but a migrated file does not exist in storage. See “Resolving Orphaned Stub Files” on page 81 for more information.

Table 14. Reconciliation Tasks (continued)

Object	Reconciliation Tasks
Premigrated files database	<p>Valid for AIX JFS only</p> <ul style="list-style-type: none"> • Deletes invalid entries from the premigrated files database. • Completes reorganization of the premigrated files database that is stored in the .SpaceMan directory if processing is interrupted. • Clears the database if the premigrated files database is corrupted. All previously premigrated files become resident files.
Status file	<p>Updates the following information in the status file:</p> <ul style="list-style-type: none"> • Number of premigrated files • Number of premigrated blocks • Number of migrated files • Number of migrated blocks. <p>Note: Migration and recall processes update status information dynamically. If any other process changes the state of a file, the status file does not reflect the change until reconciliation is run.</p>

Reconciling Your File Systems Manually

You can perform reconciliation tasks manually for one or more file systems, or synchronize your client and server only. The sections that follow describe each task. For example, if you recall a large number of migrated files, modify them, and selectively migrate them to storage, two copies of each file reside in storage. The unmodified copy of each file now is obsolete. If you set the ***migfileexpiration*** option to zero, you can run reconciliation immediately to delete the obsolete copies from storage and create available space for your migrated files.

Perform All Reconciliation Tasks

To manually perform all reconciliation tasks for one or more file systems, follow these steps:

1. Click **Space Manager** in the Hierarchical Storage Management window.
2. Highlight the file systems that you want to reconcile.
3. Click **Selected**→**Reconcile (Synchronize and Build)** to view the Reconcile window. Both check buttons for **Synchronize Client and Server** and **Build a Migration Candidates List** are selected.

If you select more than one file system, a window displays for each one. Select the window that you want to view.

4. Click **Start** to begin reconciling the file system. The Reconcile Status window displays.
5. Click **Stop** in the Reconcile Status window to discontinue file system reconciliation.

Note: The file system might not be completely synchronized, or the migration candidates list might not be completely built. If either occurs, the file system will not be reconciled until reconciliation runs again.

6. Click **Close** to close the Reconcile Status window.

After you run reconciliation, check the orphan.stubs file in the .SpaceMan directory for each file system that you reconciled to determine if any orphaned stub files were located. If the orphan.stubs file lists file names, see “Resolving Orphaned Stub Files” on page 81.

To view the migration candidates list for your file system, click the file system in the Space Manager window. Click **Selected**→**Display Migration Candidates List**.

Synchronize Your Client and Server

When you synchronize your file system, you update other space-management-related information. To synchronize your client and server manually, follow these steps:

1. Click **Space Manager** in the Hierarchical Storage Management window.
2. Highlight the file systems that you want to synchronize.
3. Click **Selected**→**Synchronize Client and Server** to view the Reconcile Status window. If you select more than one file system, the HSM client displays a window for each one. Select the window that you want to view.
4. Click **Stop** in the Reconcile Status window to discontinue synchronization.

Note: The file system might not be completely synchronized until you run **Synchronize Client and Server**, run all reconciliation tasks manually, or run reconciliation automatically.

5. Click **Close** to close the Reconcile Status window.

After you synchronize the client and server, check the orphan.stubs file in the .SpaceMan directory for each file system that you reconciled to determine if any orphaned stub files were located. If file names are listed in the orphan.stubs file, see “Resolving Orphaned Stub Files.” See Appendix B, “The .SpaceMan Directory,” on page 161 for information about the .SpaceMan directory.

Resolving Orphaned Stub Files

An orphaned stub file is a stub file for which a corresponding migrated file in storage is not located. If orphaned stub files exist in your file systems, the HSM client records information about these files in the orphan.stubs file during reconciliation. If you set the **errorprog** option in your dsm.sys file, a message is sent to the program that you specified with this option during automatic reconciliation.

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.
If files are no longer needed, an administrator can delete some or all of them from storage. In this case, the stub files are no longer valid and you can erase them.
- A media failure occurs that corrupts or loses your migrated files.
Storage pool backup and recovery provides protection against media failures. However, if you cannot restore a migrated file from a migration storage pool, you can restore a backup version of the file if you used the backup-archive client. When you set the **restoremigstate** option to *no* in your dsm.opt file and you then restore a backup version of a migrated file, the file becomes a normal, resident file.

To check for orphaned files, specify *yes* on the **checkfororphans** option in your `dsm.sys` file. When orphaned files are located, their names are recorded in the `.SpaceMan/orphan.stubs` file. If you select *yes*, a full file system tree traversal is performed. In this instance, you cannot run the **dsmautomig** command in parallel with the **dsmreconcile** command.

Reconciling Your File Systems Using Commands

You can use the **dsmreconcile** command to perform reconcile tasks rather than using the graphical user interface. If you reconcile several file systems, increase the value on the **reconcileinterval** option in your `dsm.sys` file to reduce the impact that the **dsmreconcile** command might have on system performance.

Use the **dsmreconcile** command to perform these reconciliation tasks:

- Reconcile your file systems
- Perform all other reconciliation tasks
- Synchronize your file systems.

Table 15 provides several examples of reconciliation tasks that you can perform.

Table 15. Examples Using the dsmreconcile Command

To Do This	Enter This
Reconcile the <code>/home</code> file system.	<code>dsmreconcile /home</code>
Synchronize the <code>/home</code> file system.	<code>dsmreconcile -f /home</code>

The **dsmreconcile** command traverses your managed file systems under the following conditions:

- You set the **checkfororphans** option to *yes* in your `dsm.sys` file.
- You specify **dsmreconcile -o**. This forces a full tree traversal.
- You set the **restoreMigState** option to *yes* (the default) in your `dsm.opt` file. Stub files are created during a restore.
- The HSM support is removed from your file system.

When you upgrade from a previous version of the HSM client, the next run of **dsmreconcile** on a previously managed file system forces a full tree traversal. This process occurs only once for each managed file system to update the client and server databases to the new format.

The **dsmautomig** command can run in parallel with **dsmreconcile** if it is not necessary for the **dsmreconcile** command to traverse your file system. The exception to this is when **dsmreconcile** queries the server for a list of migrated files.

For command information about **dsmreconcile**, see “`dsmreconcile`” on page 151. For command information about **dsmautomig**, see “`dsmautomig`” on page 118.

Using Space Management Daemons

The HSM client uses the space monitor daemon, the recall daemon, and the scout daemon to manage your file systems automatically. They start when you add space management to your file systems and modify your options. The sections that follow describe each space management daemon.

Start the Space Monitor Daemon

The space monitor daemon monitors space usage on all file systems to which you add space management, and it starts threshold migration whenever necessary. To check space usage more frequently or less frequently, change the value on the *checkthresholds* option in your `dsm.sys` file.

To reconcile your file systems more frequently or less frequently, change the value on the *reconcileinterval* option in your `dsm.sys` file. See Chapter 12, “Using Options,” on page 89 for more information about these options.

For AIX file systems only: The space monitor daemon starts automatically when you mount your file system and add space management to it. If the space monitor daemon stops running, enter the `dsmonitor` command to start it.

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

Start the Recall Daemon

The recall daemon recalls migrated files from storage to your local file system. A 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 recall daemons that you want to run at one time, use the *minrecalldaemons* and *maxrecalldaemons* options in your `dsm.sys` file. The minimum number of recall daemons that you can run at the same time is one. The default is three. The maximum number of recall daemons that you can run at the same time is 99. The default is 20. See “Minrecalldaemons” on page 107 and “Maxrecalldaemons” on page 101 for more information about these options.

The maximum number of recalls that you can set depends on the number of concurrent recalls that normally occur on your system. If all recall daemons are busy, another file cannot be recalled until a recall daemon is available. If a frequently-used application opens several files at the same time, and that application uses all available recall daemons because all files are migrated, increase the value that you set on the *maxrecalldaemons* option. If a recall daemon is unable to start another process that is attempting to access a migrated file, that process cannot continue until a recall daemon is available.

The master recall daemon starts automatically when you mount your file system and add space management to it. If a recall daemon is not running, enter the `dsmrecalld` command to start one.

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

Start the Scout Daemon

The scout daemon automatically searches for candidates on each file system for which space management is active. To specify how regularly a search for candidates will start on a file system, modify the setting on the *candidatesinterval* option in your `dsm.sys` file. If you modify the setting to the maximum value of 9999, the scout daemon does not search in regular cycles for migration candidates. It searches for migration candidates only on requests from the automation process.

Stop the Space Management Daemons

To stop the space monitor, master recall or subordinate recall daemons, or scout daemon, follow these steps:

1. Enter the **dsmq** command to obtain the recall ID and the recall daemon process ID for each recall process that is in the queue. For more information, see “dsmq” on page 147.
2. Enter the **dsmrm** command to remove each recall process from the queue. For more information, see “dsmrm” on page 153.
3. Enter **ps -ef | grep dsm** to verify that both the space monitor daemon and the master recall daemon are running.
4. Enter the **kill -15** command with the process identifier number to stop the daemons.

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

5. Verify that the daemons no longer are running.

Chapter 10. Using Utilities

The HSM client provides several commands that you can use to perform utilities tasks. This chapter describes these utilities.

Displaying Information About Your File Systems

Note: A graphical user interface is not available on AIX GPFS.

You can display space management-related information about each of your file systems in a list format or a pie chart format from the graphical user interface (**dsmhsm**). Click **Space Manager** in the Hierarchical Storage Management window. The Space Manager window displays.

Note: Some column headings in the File System view might appear shortened. To display the complete headings, specify a larger column width on the Tivoli Storage Manager Column Layout for Space Manager window. To open the Column Layout window, select **Column Layout** from the View menu.

In the Space Manager window, you can:

- Use the scroll bar to view information.
- Go directly to a specific column of information. Click a **Scrolling Shortcut** button at the bottom of the Space Manager window.
- Click **View**→**Column Layout** to arrange the columns of information in a different order.
- Click **View**→**Sort** to sort information according to the values in a particular column.
- Click **View**→**Refresh** or **Refresh Interval** to refresh the displayed information once or at intervals.
- Click **Selected**→**View as Pie Chart** to view information for a file system in a pie chart format. In this window, you can select options from the **View** menu to refresh information for a file system once or at intervals.
- Click **File**→**Close** to close the View as Pie Chart window.

The migration and recall processes dynamically update status information. If another process changes a file, the change is not reflected in the Space Manager or View as Pie Chart windows until reconciliation runs. For example, if you restore a backup version of a migrated file, the information in the View as Pie Chart window for local disk space, server disk space, and local disk inodes does not reflect the change until reconciliation runs. Or, if you use any command that changes the ctime of a premigrated file (causing the file to become resident), that change is not reflected in status information until reconciliation runs.

Display Information from the Command Line

You can use the commands that are listed in the following table to display space management information about your file systems, files, and directories. For information about each of these commands, see “Hierarchical Storage Management Commands” on page 114.

Table 16. Displaying Space Management Information

Command	Description
dsmdf	<p>Displays space usage information for a file system. For example, to display space usage information for the /home file system, enter:</p> <pre>dsmdf /home</pre> <p>See “dsmdf” on page 120 for more information about this command.</p>
dsmls	<p>Lists files in a directory and displays file conditions. For example, to display information about all files in the /home/user1 directory, enter:</p> <pre>dsmls /home/user1/*</pre> <p>See “dsmls” on page 122 for more information about this command.</p>
dsmdu	<p>Displays space usage information for files and directories. For example, to display space usage information for each file in the /home/user/proj1 directory and in all of its subdirectories, enter:</p> <pre>dsmdu -Allfiles /home/user1/proj1</pre> <p>See “dsmdu” on page 121 for more information about this command.</p>
dsmmigfs	<p>Displays the current space management settings for a file system. For example, to display the space management settings for the /home file system, enter:</p> <pre>dsmmigfs query /home</pre> <p>See “dsmmigfs” on page 131 for more information about this command.</p>
dsmmigundelete	<p>Recreates deleted stub files for migrated files and creates stub files for premigrated files if a corresponding original file does not exist on your local file system. The file then becomes a migrated file. For example, 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:</p> <pre>dsmmigundelete /home</pre>

Chapter 11. 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. Your administrator also sets server parameters to:

- Balance scheduled services for all client nodes.
- Specify that your client node can query the server for scheduled work at specific time intervals, or wait for the server to contact your client node when it is time to perform scheduled services.
- Control how often your client node contacts the server for scheduled work.

Central scheduling requires a cooperative effort between a Tivoli Storage Manager server and your client node.

Setting Scheduling Options

Before scheduled services can be performed, set scheduling options in your `dsm.sys` file and start a client scheduler on your workstation. The backup-archive command-line client must be installed to start the client scheduler. For information about setting these options in your `dsm.sys` file, see *IBM Tivoli Storage Manager for UNIX: Backup-Archive Clients Installation and User's Guide*.

Starting the Client Scheduler

The client scheduler must be running on your workstation before you start any scheduled work. Enter the Tivoli Storage Manager backup-archive client **schedule** command to start the client scheduler. For more information about starting the client scheduler, see *IBM Tivoli Storage Manager for UNIX: Backup-Archive Clients Installation and User's Guide*.

You can set up a cron job to run space management services at specific times. Set the **reconcileinterval** option to zero in your `dsm.sys` file so the HSM client does not reconcile file systems automatically at specific intervals. Use a cron job to run reconciliation at a specific time each day.

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.

Displaying Information About Scheduled Services

You can display information about scheduled services as well as information about completed services. To display information about scheduled services for your client node, enter the command displayed below. This command is provided with the Tivoli Storage Manager backup-archive client. For more information about the **dsmc query schedule** command, see *Tivoli Storage Manager for Unix Installing and Using the Backup-Archive Clients*.

```
dsmc query schedule
```

When you run the **dsmc query schedule** command in the foreground, output from scheduled commands displays on your screen. The output is directed also to a file named `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 to the dsm Sched.log file. Check the schedule log to verify that all work completed successfully.

Chapter 12. Using Options

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

The HSM client provides system and space management options that you can 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 they affect automatic migration, reconciliation, and recall. In your `dsm.sys` file, group your options into stanzas for each server that your client node contacts for backup, archive, and space management services.

Options are processed following this order:

1. Options that are defined on the server with server-enforced client options. The client *cannot* override the value.
2. Options that are entered locally on the command line.
3. Options that are defined on the server for a schedule using the options parameters.
4. Options that you enter locally in your options file.
5. Options that are received from the server with client options that the server does not enforce. The client can override the value.
6. Default option values.

With the exception of the ***errorprog*** option, place all space management options *before* the stanzas for each server that you define in your `dsm.sys` file. For example, if you want to specify a default server for backup and archive services, a migration server, and new values for other space management options, place the options in your `dsm.sys` file. The following is an example of a server stanza containing options for a server that you want to contact.

```
DEFAULTServer          server1
MIGRATEServer          server2
CHECKThresholds        2
CANDIDATESInterval    12
MAXCANDProcs           5
RECOncileinterval     12
MAXRECOncileproc       5
MAXThresholdproc       5
MINMIGFILESIZE        8192
MIGFILEEXPIration     10
MINRECALLdaemons      5
MAXRECALLdaemons      15
CHECKFororphans        no
MAXMIGRators           1
KERNelmessages         no
OVERLAPRECALL          no
Servername           server1
  COMMmethod            TCPip
  TCPPort               1500
  TCPServeraddress      alvmmd.almaden.ibm.com
  Passwordaccess        generate
  Mailprog              /usr/bin/xsend root
  Groups                system adsm
  Users                 steiner chron wang nguyen
  Inclexc1              /adm/adsm/backup.exc1
  ERRORProg             /bin/cat
```

Servername	server2	
COMMethod		SNAlu6.2
PARTner1uname		raptor
TPname		appcdel
CPICM0dename		appc
Passwordaccess		generate
Mailprog		/usr/bin/xsend root
Groups		system adsm
Users		sullivan tang stewart
Incl excl		/adm/adsm/migrate.excl
ERRORProg		/bin/cat

Table 17. Space Management Options in Your Options Files

Option	Description	Page
<i>candidatesinterval</i>	Specifies how often the dsmscoutd daemon searches for migration candidates in file systems to which you added space management. The range of values is one through 9999. The default is one.	92
<i>checkfororphans</i>	Specifies whether or not the ds mreconcile command checks for orphans. Specify yes or no. The default is no.	93
<i>checkthresholds</i>	Determines how frequently the space monitor daemon checks space usage on the file systems to which you add space management. The range of values is one through 9999. The default is five.	94
<i>compression</i>	Compresses files before you send them to the Tivoli Storage Manager server. Specify yes or no. The default is no.	95
<i>defaultserver</i>	Specifies the name of the default Tivoli Storage Manager server to contact for space management services if you do not specify a server name on the migrateserver option. You cannot override the default server for space management services in your dsm.opt file.	96
<i>errorprog</i>	Specifies a program to which you want to send a message if a severe error occurs during space management processing. Place this option <i>after</i> all space management options for each server that you define in your dsm.sys file.	97
<i>kernelmessages</i>	Valid for AIX JFS file systems only Displays HSM-related messages during processing. If you change the value, it will not take effect until you restart your system. Specify yes or no. The default is yes.	98
<i>maxcandprocs</i>	Specifies the number of slave scout daemons that can scan for migration candidates in parallel on a system. The range of values is two through 20. The default is five.	99
<i>maxmigrators</i>	Specifies the number of parallel migration sessions that you can perform. The range of values is one through 20. The default is one.	100
<i>maxrecalldaemons</i>	Specifies the maximum number of recall daemons that you can run at one time. The range of values is two through 99. The default is 20.	101
<i>maxreconcileproc</i>	Specifies the maximum number of automatic reconciliation processes that you can start at one time. The range of values is one through 99. The default is three.	102
<i>maxthresholdproc</i>	Specifies the maximum number of automatic threshold migration processes that you can start at one time. The range of values is one through 99. The default is three.	103
<i>migfileexpiration</i>	Specifies the number of days that copies of migrated or premigrated files remain on the Tivoli Storage Manager server after they are modified on, or erased from, your local file system. The range of values is zero through 9999. The default is seven days.	104

Table 17. Space Management Options in Your Options Files (continued)

Option	Description	Page
migrateserver	Specifies the Tivoli Storage Manager server to which you want to migrate files from your client node. Specify only one migration server for each client node. You cannot override the migration server in your dsm.opt file.	105
minmigfilesize	Specifies the minimum file size for a file to be eligible for automatic migration. The range of values is the fragment size of the file system through 2147483647. The default is zero.	106
minrecalldaemons	Specifies the minimum number of recall daemons that you want to run at the same time to perform recalls for your client node. The range of values is one through 99. The default is three.	107
optionformat	Specifies the format to use (either standard or short) when you specify HSM client commands. If you change the value, the new value is effective immediately. Select standard or short format. The default is standard.	108
overlaprecall	Valid for AIX JFS file systems only Specifies whether migrated files must be recalled immediately even if a process is accessing only the resident part of the file. This option is valid for read-only operations on a file. Specify Yes or No. The default is No.	109
reconcileinterval	Specifies how often automatic reconciliation of file systems occurs to which you add space management. The range of values is zero through 9999. The default is 24.	110
restoremigstate	Restores a file if it is backed up after migration. You cannot restore a stub file if you back up a file before migration because a server stub file copy does not exist. Specify yes or no. The default is yes.	111

Candidatesinterval

The ***candidatesinterval*** option specifies how often the dsmscoutd daemon searches for migration candidates in file systems to which you added space management.

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

Syntax

▶▶—Candidatesinterval *number_in_hours*————▶▶

Parameters

number_in_hours

Specifies the maximum number of hours that will elapse between each successive time the dsmscoutd daemon automatically searches for candidates in a file system. The range of values is one through 9999; the default is one. If it is required, the dsmscoutd daemon will scan a file system more often than you specify.

If you modify the setting to the maximum value of 9999, the dsmscoutd daemon does not search in regular cycles for migration candidates. It searches for migration candidates only on requests from the automigration process.

Checkthresholds

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

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

Syntax

▶▶—CHECKThresholds *interval*—————▶▶

Parameters

interval

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

Compression

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

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. 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.

This option controls compression *only if* your administrator specifies that your client node determines the selection.

Note: The server also can define this option.

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 also can specify the server to which files are migrated from your local file systems.

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

Syntax

▶▶—DEFAULTServer *servername*—————▶▶

Parameters

servername

Specifies the name of the default server to which you back up and archive your files. Specify the server to which your files migrate from your local file systems. Use this option if you do not specify a migration server with the ***migrateserver*** option.

Errorprog

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

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

Syntax

▶—`ERRORProg program-name`—▶

Parameters

program-name

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

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

Kernelmessages

Valid for AIX JFS file systems only

The ***kernelmessages*** option specifies whether or not to display HSM-related messages during processing.

Note: During transparent recall, if more than one process accesses the same migrated file at one time, the kernel displays a message only for the first process that accesses the file.

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

Syntax



Parameters

Yes

Displays HSM-related messages that the kernel sends. This is the default.

No

Does not display HSM-related messages that the kernel sends.

Maxcandprocs

The *maxcandprocs* option specifies the number of slave scout daemons that can scan for migration candidates in parallel on a system.

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

Syntax

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

Parameters

max_number_of_scout_daemons

Specifies the maximum number of slave scout daemons that you can set. The range of values is two through 20; the default is five.

Maxmigrators

The ***maxmigrators*** option specifies the maximum number of parallel migration sessions per file system to the Tivoli Storage Manager server that the ***dsmautomig*** command can perform. Ensure that you have sufficient resources on the server for parallel migration to occur. Do not set the ***maxmigrators*** option higher than the number of parallel sessions that the server can use to store data.

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

Syntax

▶▶—MAXMIGRrators *number*—————▶▶

Parameters

number

Specifies the maximum number of parallel migration sessions that you can set. The range is one through 20; the default is one.

Maxrecalldaemons

The *maxrecalldaemons* option specifies the maximum number of recall daemons that you can run at one time to perform recalls for your client node. During normal operations, if the number of recall daemons that are running at one time is similar to the maximum number that is permitted, increase the value. For example, if you use an application that opens several files at one time, and these files are migrated, the application can use all available recall daemons. If you cannot start a recall daemon for another process that is attempting to access a migrated file, the process stops until a recall daemon is available.

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

Syntax

▶▶—MAXRECALLdaemons *number*—————▶▶

Parameters

number

Specifies the maximum number of recall daemons that you can run at one time to perform recalls. The range is two through 99; the default is 20.

Maxreconcileproc

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

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

Syntax

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

Parameters

maxreconcileproc

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

Maxthresholdproc

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

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

Syntax

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

Parameters

maxthresholdproc

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

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 deleted from your local file system.

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

Syntax

▶▶—MIGFileexpiration *days*—————▶▶

Parameters

days

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

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. You cannot override the server that you specify for migration in your `dsm.opt` file, or when you enter a command.

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.

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

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

Syntax

▶▶—MIGRateserver *servername*—————▶▶

Parameters

servername

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

Minmigfilesize

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

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

Syntax

▶▶—MINMIGfilesize *fileSize*—————▶▶

Parameters

fileSize

Specifies the minimum file size, in bytes, for a file to be eligible for migration. This byte size is ignored if it is smaller than the file system block size. The range of values is zero through 2147483647; the default is zero.

If you specify the default, the HSM client uses the file system block/fragment size or stubsize as the minimum size for files that can be migrated; whichever is larger.

Minrecalldaemons

The *minrecalldaemons* option specifies the minimum number of recall daemons that you can run at the same time to perform recalls for your client node.

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

Syntax

▶▶—MINRecalldaemons *number*—————▶▶

Parameters

number

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

Optionformat

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

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

Syntax

►—OPTIONFormat—Standard
Short—►

Parameters

Standard

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

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

Short

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

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

Overlaprecall

Valid for AIX JFS file systems only

The ***overlaprecall*** option specifies whether migrated files must be recalled immediately even if a process is accessing only the resident part of the file. This option is valid for read-only operations on a file.

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

Syntax



Parameters

No Migrated files will not be recalled immediately. This is the default.

Yes

Migrated files will be recalled immediately.

Reconcileinterval

The *reconcileinterval* option specifies how often the space monitor daemon reconciles your file systems. This includes marking deleted local files for expiration on the server, removing stale entries from the premigration database, and bringing the status file up to date.

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

Syntax

►►—RECOncileinterval *interval*—————►►

Parameters

interval

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

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 13. Using Commands

You can use hierarchical storage management (HSM) commands rather than the graphical user interface (GUI) to perform most space management tasks. This chapter provides information about using these commands. For information about client system options, see Chapter 12, “Using Options,” on page 89.

Note: The GUI is available for root users only. It is not available for HSM on AIX GPFS or HP-UX VxFS.

Entering Commands

When you enter commands, follow these rules:

- *Do not* precede HSM commands with **dsmc**. Each HSM client command is a separately-executable command.
- Enter the complete command name in lowercase letters. You cannot use uppercase letters or an abbreviation for a command name.
- Use the following wildcard characters in file, directory, or file system specifications. The shell in which you are running matches and expands wildcard characters.
 - * Matches zero or more characters
 - ? Matches any single character
- Enter the characters in a command in a continuous string without pressing the Return key. You can enter as many as 256 characters on the command line.

Enter Options With Commands

Use either a standard format (the complete option name), or a short format (a one or two-letter form of the option name) to enter options. The standard format is the default. Specify the format (either standard or short) on the *optionformat* option in your `dsm.opt` file. If you do not specify a format, the default is used.

Use the Standard Format

If you set the *optionformat* option to standard in your `dsm.opt` file, enter options using the complete option name. For example, enter commands and options in this manner:

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

Follow these rules when you use options with commands:

- Enter options in any combination of uppercase and lowercase letters. Options are not case-sensitive.
- Precede each option with a hyphen (-). For example:

```
dsmmigquery -candidatelist -premigratedlist /home
```
- Separate each option or parameter with a blank space.
- Enter more than one option in a command in any order before or after a file, directory, or file system specification. For example, you can enter either of the following:

```
dsmmigrate -recursive -detail /home/user1/test1  
dsmmigrate /home/user1/test1 -detail -recursive
```

Use the Short Format

If you set the *optionformat* option to short in your `dsm.opt` file, enter options using an abbreviation. Uppercase letters in each option description indicate the minimum abbreviation that is permitted. For example, enter commands and options in this manner:

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

Follow these rules when you use options with commands:

- Enter options exactly as they display in the help for each command. Options are case-sensitive.
- Enter the options that you want to use in a command in any order *before* you enter a file, directory, or file system specification. For example, you can enter either of the following:

```
dsmmigrate -R -v /home/user1/test1  
dsmmigrate -v -R /home/user1/test1
```

- Group more than one option, except any that require an argument, in any order either after one hyphen (-), or with a hyphen (-) for each option. For example, you can enter either of the following:

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

- Separate each option or parameter with a blank space if you do not group options after one hyphen (-).

Displaying Online Help for Commands

You can display online help for HSM commands in either of the following ways:

- Enter the *-help* or *-h* option with any command. For example:

```
dsmmigrate -help  
dsmmigrate -h
```

The command syntax displays.

- Enter the **dsmmighelp** command. A list of help topics displays from which you can select general help information for commands, or help for a specific command.

Note: Because the client messages and online glossary are included with the UNIX backup-archive command-line help, they are not repeated with the HSM client. To access client messages or the glossary, access them through the Tivoli Storage Manager for UNIX backup-archive client program.

Hierarchical Storage Management Commands

Table 18 provides an alphabetical list of the HSM client commands, a brief description of each command, and the command page number.

Table 18. HSM Commands

Command	Description	Page
dsmattr	Valid for AIX JFS and AIX GPFS only Sets or displays the recall mode for a migrated file.	116
dsmautomig	Starts parallel migration sessions for a file system.	118
dsmdf	Displays space usage information for a file system.	120

Table 18. HSM Commands (continued)

Command	Description	Page
dsmdu	Displays space usage information for files and directories.	121
dsmls	Lists files in a directory and displays file state.	122
dsmmigfs	Adds space management to a file system, or updates space management attributes for a file system.	124
	Deactivates or reactivates space management for a file system, or removes space management from a file system.	127
	Manages recovery from partial system failure (GPFS only).	130
	Displays current space management settings for a file system.	131
	Deactivates or reactivates space management for your client node.	133
	Imports and exports the HSM management of a file system within an HACMP environment.	128
	Transfers the HSM management of a file system to a local node within the same local GPFS node set.	134
	Transfers the HSM management of a file system to the preferred node if the node is different from the current owner node.	135
dsmmighelp	Displays online help for commands.	136
dsmmigquery	Displays space management information.	137
dsmmigrate	Moves selected files from your local file system to Tivoli Storage Manager storage.	139
dsmmigundelete	Recreates deleted stub files.	141
dsmmode	Valid on AIX JFS file systems only	143
	Sets one or more execution modes that affect the space management-related behavior of commands.	
dsmmonitor	Starts the space monitor daemon.	146
dsmq	Displays information, including recall IDs, for all files that are currently queued for recall.	147
dsmrecall	Moves selected files from storage to your local file system.	148
dsmrecalld	Starts the recall daemon.	150
dsmreconcile	Synchronizes the client and server.	151
dsmrm	Removes a recall process from the recall queue.	153
dsmrootd	Starts the root daemon.	154
dsmscoutd	Valid on AIX JFS, Solaris VxFs, and HP-UX VxFs file systems only	155
	Starts the scout daemon.	
dsmsetpw	Changes the Tivoli Storage Manager password for your client node.	156
dsmwatchd	Valid on AIX GPFS and Linux86 GPFS file systems only	157
	Manages failover activities for your local node.	

Value	Description
<i>Partialrecall (p)</i> (AIX GPFS HSM only)	Specifies that the file should be recalled using partial file recall, regardless of its size.
<i>Streaming (s)</i> (AIX GPFS HSM only)	Specifies that you want to enable an asynchronous recall of migrated files. The recalled portion of the file can be accessed while the file is recalled. This parameter is valid for read-only operations on the file.

-RECURsive or -R

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

filespec

Attention: This parameter is required only when you set a new recall mode. If you do not use the **-RECALLmode** option, and you do not specify a path and a file name, the current recall mode displays for all files in the current directory.

The path and file name of the file for which you want to set a new recall mode, or display the current recall mode. You can specify a single file, a group of files, or a directory. If you specify a directory, the HSM client sets or displays that recall mode for each migrated file in the directory.

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

Examples

Task	Command
Change the recall mode to migrate-on-close for a file named /home/user1/file1.	<code>dsmattr -recall=migonclose /home/user1/file1</code>
Change the recall mode to read-without-recall for all migrated files that begin with JAN in the /home/user1 directory.	<code>dsmattr -recall=readwithoutrecall /home/user1/JAN*</code>
Change the recall mode to migrate-on-close for all migrated files in the /home/user1 directory and all of its subdirectories.	<code>dsmattr -recall=migonclose -R /home/user1</code>
Change the recall mode to partial file recall for all migrated files in the /home/user2 directory and all of its subdirectories.	<code>dsmattr -recall=partialrecall -R /home/user2</code>
Enable an asynchronous recall of migrated files in the /home/user2/ directory.	<code>dsmattr -recall=streaming /home/user2/</code>
Display the recall modes that are assigned to all files in the current directory.	<code>dsmattr</code>

dsmautomig

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 Tivoli Storage Manager server.
- Checks if a current backup version exists.
- Starts parallel migration sessions to the Tivoli Storage Manager server, migrating more than one file at a time.

Note: If the LANG environment variable is set to C, POSIX (limiting the valid characters to those with ASCII codes less than 128), or other values with limitations for valid characters, the HSM client skips files which have file names containing invalid characters with ASCII codes higher than 127. If you are using a single-byte character set (SBCS) such as English as your language environment, all file names are valid and are migrated by the HSM client.

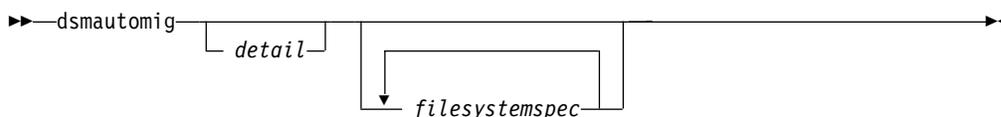
Multi-byte characters are interpreted as a set of single bytes all containing valid characters. If you are using multi-byte character sets (MBCS) as your language environment, the HSM client migrates file names that consist of valid characters in the current environment. For example, a file name consisting of Japanese characters may contain invalid multi-byte characters if the current language environment is a Chinese character set. Files containing invalid multi-byte characters are not migrated or recalled. If such files are found during migrate or recall no information is printed. The HSM daemons must run in the en_US language locale (or another SBCS language locale) to work properly.

Specify the number of parallel migration sessions with the **maxmigrators** option in your dsm.sys file. Verify that sufficient resources are available on the Tivoli Storage Manager server for parallel migration. Do not set the **maxmigrators** option higher than the number of sessions that the Tivoli Storage Manager server can use to store data. Start threshold migration manually to lower space usage on your file system *before* it reaches the high threshold that you set.

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.

Note: The **dsmautomig** and **dsmreconcile** processes must be found with the PATH variable, or the dsmonitord daemon cannot perform reconciliation and threshold migration.

Syntax



Parameters

- detail or -v
Displays information about migrated files.

filesystemspec

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.

Examples

Task	Command
Start threshold migration for all file systems for which space management is active.	<code>dsautomig</code>
Start threshold migration for the <code>/home</code> file system.	<code>dsautomig /home</code>
Start threshold migration for the <code>/home</code> and <code>/test1</code> file systems.	<code>dsautomig /home /test1</code>

dsmdf

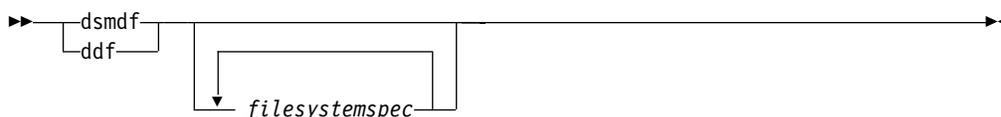
The **dsmdf** command displays the following information for one or more file systems:

- File system state: active (a), inactive (i), global inactive (gi), or native (n).
For AIX GPFS and Linux86 GPFS file systems only: Remotely-managed file systems within the local GPFS node set (r)
- 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 Linux86 GPFS file systems only: Migrating newly-created files of less than five minutes might display incorrect results (resident size) when you use the **dsmdf** and **dsmdu** commands. This is because GPFS is not synchronized on all nodes when you migrate files. The last block of a file is not released from the disk although the file migrated successfully. This can cause a deviation from an assumed disk usage if many small files are migrated and the block size is high.

Syntax



Parameters

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

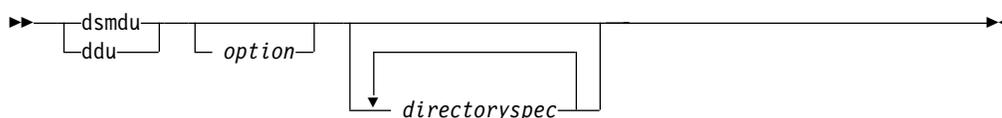
Task	Command
Display information for all file systems to which you added space management.	dsmdf
Display information for the /home file system.	dsmdf /home

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 Linux86 GPFS file systems only: Migrating newly-created files of less than five minutes might display incorrect results (resident size) when you use the **dsmdf** and **dsmdu** commands. This is because GPFS is not synchronized on all nodes when you migrate files. The last block of a file is not released from the disk although the file migrated successfully. This can cause deviation from assumed disk usage if many small files are migrated and blocksize is high.

Syntax



Parameters

options

If you do not specify either of the following options, the HSM client displays the number of 1 KB blocks that the specified directory and each of its subdirectories use.

-Allfiles or **-a**

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

-Summary or **-s**

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

directoryspec

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

Examples

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

dsmls

The **dsmls** command displays the following information about a list of files:

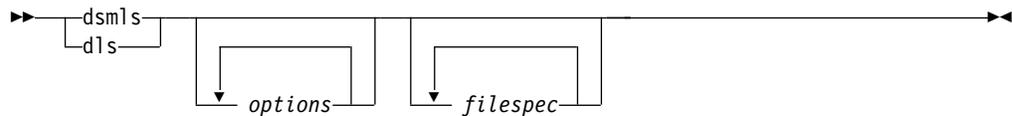
- Actual size (in bytes)
- Resident size (in bytes)
- Resident block size (in KB)
- File state and recall mode
- File name.

For a resident or premigrated file, the actual size and resident size are the same. For a migrated file, the actual size is the size of the original file. The resident size is the size of the stub file that remains on your local file system.

The file state for a file can be any of the following: migrated (m), premigrated (p), or resident (r). A dash (-) indicates a directory or a non-regular file; for example, a character special file or a named pipe file. For a migrated file, the **dsmls** command also indicates the recall mode that you set for the file:

- If you set the recall mode to normal, additional information does not appear in the File State column.
- If you set the recall mode to migrate-on-close, the notation (m) displays in the File State column.
- If you set the recall mode to read-without-recall, the notation (r) displays 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.

Syntax



Parameters

options

Use any of the following options:

-Noheader or **-n**

Omits column headings from the output for this command.

-Recursive or **-R**

Displays information about files in subdirectories of the directory.

filespec

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

Examples

Task	Command
List all files in the current directory.	<code>dsmls</code>
List all files in the <code>/migfs2/test</code> directory.	<code>dsmls /migfs2/test</code>
List all files in the <code>/migfs2/test</code> directory and in its subdirectories.	<code>dsmls -Recursive /migfs2/test</code>
List all files whose names begin with tf in the <code>/migfs2/test</code> directory.	<code>dsmls /migfs2/test/tf*</code>

dsmmigfs

You must have root user authority to use this command.

The **dsmmigfs** command adds space management to your file system, or updates space management settings for your file system.

For AIX GPFS and Linux86 GPFS file systems only: Before you run **dsmmigfs add filesystemName**, ensure that file system is mounted and enabled for DMAPI management. Run the following:

For AIX GPFS:
`/usr/lpp/mmfs/bin/mmlsfs DevicePath`

For Linux86 GPFS:
`/opt/lpp/mmfs/bin/mmlsfs DevicePath`

Search for the status of the **-z** flag. If DMAPI is not enabled on your file system, run the following:

For AIX GPFS:
`/usr/lpp/mmfs/bin/mmchfs DevicePath -z yes`

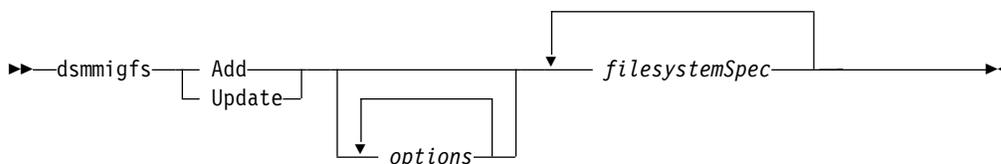
For Linux86 GPFS:
`/opt/lpp/mmfs/bin/mmchfs DevicePath -z yes`

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

In an HACMP or IBM Cluster 1350 environment: In the event of a system failure, specify the *import* parameter to access your data. The result permits an individual file system to specify a certain server stanza to manage your file system. Your file system can contain migrated files but it is not required to have them. The connection is made by adding an entry for the file system to `/etc/tsm/SpaceMan/config/dsmmigfstab`. A link is also created for the status file in `/etc/tsm/SpaceMan/status`. The *export* parameter removes the connection.

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

Syntax



Parameters

Add Adds space management to your file systems.

Update

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

If you change the high and low thresholds or the premigration percentage, the new values take effect immediately. If you change the stub file size, the

new size is used only for files that are migrated after you make the change. If you change the age or size factor, the new value is used the next time a migration candidates list is built.

If you reduce the quota, and the amount of migrated and premigrated data exceeds the new quota, the HSM client does not migrate any additional files until enough files are recalled during automatic recall or selective recall to drop the total number of megabytes for migrated and premigrated files below the new quota.

filesystemSpec

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

options Use the option settings that are provided for this command to add or update space management settings. See the table that follows for options and information.

Table 19. Space Management Option Settings

Option	Short Name	Description
-HThreshold= <i>n</i>	- <i>tn</i>	The high threshold percentage that you set for space usage on your file systems. Specify a value of zero through 100 percent. The default is 90 percent.
-Lthreshold= <i>n</i>	- <i>ln</i>	The low threshold percentage that you set for space usage on your file systems. Specify a value of zero through 100 percent. The default is 80 percent.
-Maxcandidates= <i>n</i>	- <i>mn</i>	The maximum number of migration candidates the dsmscoutd daemon searches for during one scan period. When this number is reached, dsmscoutd stops scanning. The range of values that you can enter are nine through 9999999. The default is 10000.
-MINPartialrecallsize= <i>n</i> (AIX GPFS HSM only)	- <i>nn</i>	Specifies the minimum size (in megabytes) that a file must have to qualify for partial file recall. The range of values is 0 to 999999999. A value of 0 disables partial file recall for all files; this is the default.
-MINStreamfilesize= <i>n</i> (Valid for AIX GPFS, Linux86 GPFS, and Solaris VxFs)	- <i>nn</i>	Specifies a number to enable or disable an asynchronous recall of migrated files. This number is the number of bytes that must be recalled before HSM starts streaming data to the requesting application (to ensure a steady stream of data). The recalled portion of the file can be accessed while the file is recalled. The range of values is zero through 2147483647. The default is zero. Zero disables the asynchronous option.
-Pmpercentage= <i>n</i>	- <i>pn</i>	The percentage of file system space that is available to contain premigrated files. The default is the difference between the percentage that you set for the high threshold and the percentage that you set for the low threshold. Specify a value from zero through 100 percent.
-Agefactor= <i>n</i>	- <i>an</i>	Assign an age factor to your files in each file system to which you added space management. The age and size factors determine the order in which eligible files migrate to Tivoli Storage Manager storage. Specify a value of zero through 999999999. The default is one.

dsmmigfs – adding or updating

Table 19. Space Management Option Settings (continued)

Option	Short Name	Description
-Sizefactor= <i>n</i>	- <i>in</i>	Assign a size factor to your files in each file system to which you added space management. The age and size factors determine the order in which eligible files migrate to Tivoli Storage Manager storage. Specify a value of zero through 999999999. The default is one.
-Quota= <i>n</i>	- <i>qn</i>	The maximum number of megabytes of data that you can migrate and premigrate from your file system to Tivoli Storage Manager storage. Specify a value from zero to 999999999. The default is the number of megabytes that are allocated for your file system.
-STubsize= <i>n</i>	- <i>stn</i>	The size of stub files remaining on your local file systems when files migrate to Tivoli Storage Manager storage. <ul style="list-style-type: none">• For AIX GPFS and Linux86 GPFS file systems: Specify a value of zero or a multiple of the physical block size. The default is zero.• For AIX JFS file systems: Specify a value of <i>size</i> fragment size or a multiple of the physical block size. The default is the fragment size.• For Solaris VxFS and HP-UX VxFS file systems: Specify a value of one or a multiple of the physical block size. The default is the block size.
-SServer=< <i>servername</i> >	- <i>Sserver name</i>	Overrides the default migration server for this file system. Specify the server to contact for space management services. Define the server in a stanza in your dsm.sys file. If you do not specify a server name, the default migration server is used. Use a dash (-) to set the server to the default migration server. After you use the <i>import</i> parameter, perform a reconcile to update the status file.

Examples

Task	Command
Add space management to the /home file system and set the high threshold to 80 percent and the low threshold to 70 percent.	<code>dsmmigfs Add /home -t80 -l70</code>
Add space management to more than one file system and accept the default values for all space management settings.	<code>dsmmigfs Add /home /test1 /proj*</code>
Update the space management settings for the /home file system as follows: <ul style="list-style-type: none">• Change the high threshold to 80 percent.• Change the low threshold to 70 percent.• Set the size factor to zero to ensure that the number of days since a file was last accessed is the only factor that determines when eligible files are prioritized for migration.	<code>dsmmigfs Update -t80 -l70 -i0 /home</code>
Specify the minimum size of files in the /home/user1 file system that will be recalled using partial file recall.	<code>dsmmigfs Update -minp=100 /home/user1</code>

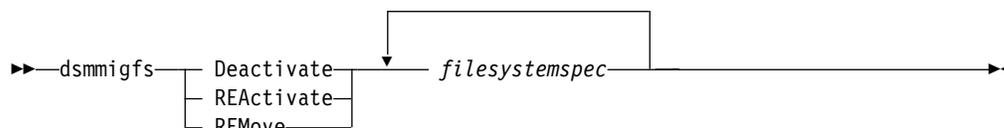
dsmmigfs

You must have root user authority to use this command.

Use the **dsmmigfs** command with the *deactivate*, *reactivate*, *remove* or *export* parameter to do the following:

- Deactivate space management for a file system.
- Reactivate space management for a file system.
- Remove space management from a file system.

Syntax



Parameters

Deactivate

Deactivates space management for a file system. The HSM client cannot perform migration, recall, or reconciliation for the file system. However, you can update space management settings for your file system, and access resident and premigrated files.

REActivate

Reactivates space management for a file system.

REMove

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

filesystemspec

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

Examples

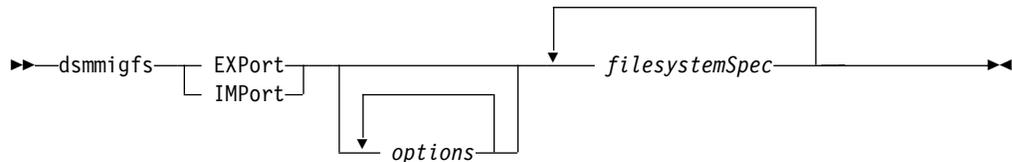
Task	Command
Deactivate space management for the /home file system.	<code>dsmmigfs Deactivate /home</code>
Reactivate space management for the /home file system.	<code>dsmmigfs REActivate /home</code>
Remove space management from the /home file system.	<code>dsmmigfs REMove /home</code>

dsmmigfs

Valid on AIX JFS file systems for HACMP and Linux86 GPFS file systems for IBM Cluster 1350 environments only. You must have root user authority to use this command.

The **dsmmigfs** command allows an HSM-managed file system to be part of a resource group in an HACMP or IBM Cluster 1350 cluster. In the event of a system failure, you can obtain access to the data from another system.

Syntax



Parameters

Export

Removes an HSM-managed file system from the current file system without recalling the file data. When you export space management from a file system, HSM removes the entry for your file system `/etc/adsm/SpaceMan/config/dsmmigfstab` and unmounts the fsm-mounted file system. Enter this command before the AIX or Linux86 volume group export (`exportvg`).

Note: The **export** command does not reconcile the file system or delete the `.SpaceMan` directory.

Import

When you import a file system, HSM connects a takeover file system to your local system. Use this command after the AIX or Linux86 import volume group (`importvg`) occurs. Run the **dsmreconcile** command to update the status file for your file system after you enter **dsmmigfs import**.

-Server=server name option

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 is defined in your `dsm.opt` file is used. To set the server to the default migration server, you also can specify the string `"-"` for the server option.

Note: This option applies to the import parameter only.

filesystemspec

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

Examples

Task	Command
Export space management for the /home file system.	<code>dsmmigfs export /home</code>
Import the space-managed file system, /home, for server, HSMServer.	<code>dsmmigfs import /home -server=HSMServer</code>

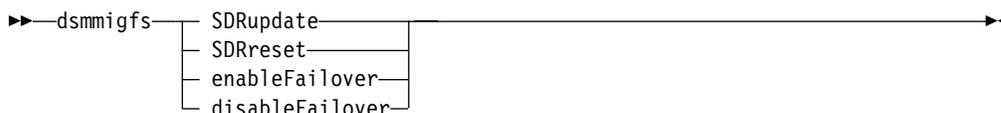
dsmmigfs

Valid on AIX GPFS and Linux86 GPFS file systems only. You must have root user authority to use this command.

The **dsmmigfs** command used with the *SDRreset*, *SDRupdate*, *enableFailover*, or *disableFailover* parameter manages recovery from partial system failure. One HSM client can take over from an HSM client that is involved in a partial system failure if the following conditions are true:

- There is one or more backup AIX or Linux86 cluster nodes within the same GPFS node set.
- The HSM client is installed on all of the backup nodes.
- The space-managed file system is mounted on at least one of the backup nodes.
- The HSM client participates in failure and takeover activities for the GPFS node set on the backup nodes where the file system is mounted. Failover has been enabled on the HSM client node and on at least one backup node.

Syntax



Parameters

Note: These SDR-related parameters also are valid for an AIX cluster (IBM Regatta pSeries 690) even if a System Data Repository is not available on those clusters.

SDRupdate

The *dsm.opt* file and the *dsm.sys* file are stored in the SDR. The set of two options files in the SDR become the single, shared set of options files for all HSM clients within the local GPFS node set. After you update your *dsm.opt* or *dsm.sys* file, use this command to synchronize all participating nodes within a GPFS node set.

SDRreset

Do not use this command during normal operation.

The HSM client attempts to reset potential locking problems in the SDR. If a command or a failover operation ended abnormally, this command will help to achieve a consistent system state.

enableFailover

Activates the local node for failover operations within the local GPFS node set.

disableFailover

Deactivates failover operations on the local node.

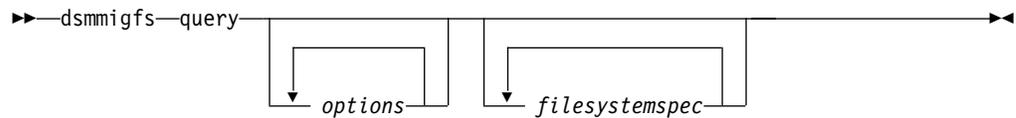
Examples

Task	Command
Deactivate failover operations on the local node.	<code>dsmmigfs disableFailover</code>

dsmmigfs

The **dsmmigfs** command used with the *query* parameter displays the current space management settings for a file system.

Syntax



Parameters

Query or -q

Displays the current space management settings for the named file system. The settings include values for the following:

- File system name
- High and low thresholds
- Premigration percentage
- Age and size factors
- Quota
- Stub file size
- Server name
- Maximum candidates
- Partial file recall
- Streaming recall
- Server stanza

Note: These settings are only displayed if you specify *-Detail* or *-v*: *maxcandidates*, *minpartialrecallsize*, and *minstreamfilesize*.

options

-Detail or -v

Displays HSM settings for each file system. Displays options that the HSM client does not display in the standard format, such as the maximum number of candidates that are located during one reconciliation.

For AIX GPFS and Linux86 GPFS file systems only: The **dsmmigfs -query** command displays only locally-managed file systems. The **dsmmigfs q -d** command displays information for all space-managed file systems within the local GPFS node set. The HSM client displays the following additional information:

- Node name for each node ID and frame ID for both the owner and the preferred nodes.
- Preferred node ID. The preferred node defines the node where HSM was initially added to the filesystem.
- Frame ID.
- Current owner node ID. The owner node ID defines the node that is currently managing the filesystem.
- GPFS node set ID.

The normal **dsmmigfs query** will not provide any GPFS-specific information (no node set ID).

dsmmigfs - Displaying Space Management Settings

Note: The failover environment is active on the local node. The **dsmmigfs query -detail** displays the current status of the local failover environment (either active or inactive).

-failover or -f

Provides a status overview of the failover environment of all HSM-managed cluster nodes. The output displays the status for the node name, node ID and frame ID. It can be any of the following:

- Active. The node participates in the failover environment within the local GPFS node set.
- Deactivated by User. You disabled failover using the **dsmmigfs disableFailover** command.
- Deactivated by HSM. The HSM client disabled failover because of an unrecoverable condition.

filesystemspec

The file system name that displays current space management settings. The default is all file systems to which space management has been added.

Examples

Task	Command
Display the current space management settings for the /migfs2 file system.	<code>dsmmigfs query /migfs2</code>

dsmmigfs

You must have root user authority to use this command.

Use the **dsmmigfs** command with the *Globaldeactivate* or *Globalreactivate* parameter to:

- Deactivate space management for the HSM-managed client node.
- Reactivate space management for the HSM-managed client node.

Syntax



Parameters

GLOBALDeactivate

Deactivates space management for all file systems on your client node. The HSM client cannot perform migration, recall, or reconciliation for any file system. However, you can update space management settings for file systems, add space management to additional file systems, or access resident and premigrated files.

GLOBALReactivate

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

Examples

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

dsmmigfs

Valid on AIX GPFS and Linux86 GPFS file systems only. You must have root user authority to use this command.

The **dsmmigfs** command transfers the HSM management of filesystem fileSpec to the local node within the same local GPFS node set.

Syntax

▶▶—dsmmigfs takeover— *filespec*—————▶▶

Parameters

takeover

Transfers the HSM management of filesystem filespec to the local node within the same local GPFS node set.

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.

Examples

dsmmigfs

Valid on AIX GPFS and Linux86 GPFS file systems only. You must have root user authority to use this command.

The **dsmmigfs** command transfers the HSM management of file systems to the preferred node if the node is different from the current owner node.

Syntax

▶▶—dsmmigfs rollback—————▶▶

Parameters

rollback

Transfers the HSM management of file systems to the preferred node if the node is different from the current owner node. Enter this command on the preferred node.

Examples

Task	Command
Transfer to the preferred node.	dsmmigfs rollback

dsmmighelp

The **dsmmighelp** command displays online help topics from which you can select general help for commands or message information.

Syntax

▶▶—dsmmighelp—————▶▶

Examples

Task	Command
Display online help for HSM commands.	dsmmighelp

dsmmigquery

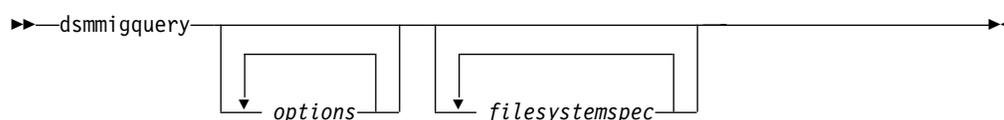
You must have root user authority to use this command.

The **dsmmigquery** command displays the following information for one or more file systems:

- Migration candidates list
- Ordered recall list for migrated files
- Available management classes
- Current client and server options
- List of all files in the file system.

Output from this command is directed to **stdout**. Use redirection characters and a file name at the end of the command to redirect the output to a file.

Syntax



Parameters

Options

Select any of these options:

-Candidatelist or **-c**

Displays the current prioritized list of migration candidates for your file system. This is the default option.

-SORTEDMigrated or **-m**

Lists all files that you migrated from the file system to Tivoli Storage Manager storage in the most efficient order for recall.

-SORTEDAll or **-s**

Lists all files in the file system in this order: resident files, premigrated files, migrated files. Sorts migrated files in the most efficient order for recall.

-Mgmtclass or **-g**

Displays information about each management class that you can assign to your files.

-Detail or **-v**

Use with the **-Mgmtclass** option to display information about each available management class. If you do not use this option, the HSM client displays the management class name and a brief description only.

-Options or **-o**

Displays the current settings for your client and server options.

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 a migration candidates list for the /migfs1 file system.	<code>dsmmigquery /migfs1</code>
Display information about management classes that you can assign to files on your client node.	<code>dsmmigquery -Mgmtclass -Detail</code>

dsmmigrate

Attention: On large file systems, selective migration can take a while.

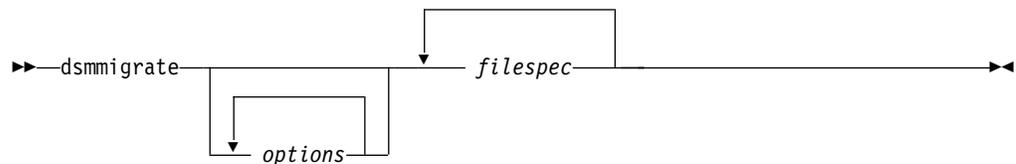
The **dsmmigrate** command selects specific files from your local file system and migrates them to a Tivoli Storage Manager server.

If the file resides in a file system with a different server stanza than the last processed file, a new session starts for each file that is migrated. This can happen as a result of links from one file system to another.

Note: If you set the *tapeprompt* option to *yes* in your *dsm.opt* file, and the destination for your migrated file is a storage pool that consists of removable media such as tape, you are prompted either to wait for the medium to mount, or to skip the file.

The first file migrates even if the file size exceeds the quota that you specified for the file system. When migration occurs, the **ddf** command displays zero migrated and premigrated bytes for your file system. If the total number of bytes exceeds the quota after the file migrates, the next file is not migrated.

Syntax



Parameters

Options

Select any of these options:

-Recursive or **-r**

Migrates files in any subdirectory below the specified directory that match the file specification. If you do not use this option, only those files from the directories that you specify are migrated.

-Detail or **-v**

Displays the size and file name for each file that you migrate.

filespec

The path and file name of the file that you want to migrate. This parameter is required. You can use wildcard characters to specify a group of files or all files in a directory. Or, you can enter more than one file specification in one command. If you enter more than one file specification, separate each specification with one or more blank spaces.

Examples

Task	Command
Migrate all files in a directory named /migfs2/test/dir1 and in all of its subdirectories. Display the information.	<code>dsmmigrate -Rv /migfs2/test/dir1</code>

dsmmigrate

Task	Command
Migrate a file named tf04 from the current directory and display the information.	<code>dsmmigrate -Detail tf04</code>

dsmmigundelete

You must have root user authority to use this command.

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.

When you delete a stub file or an original copy of a premigrated file from your local file system, the corresponding migrated or premigrated file is marked for expiration when reconciliation runs again.

Note: The **dsmmigundelete** command does not support hardlinked files. If you attempt to recreate a stub file for a hardlinked file, a stub file is not recreated *unless* all of the files that are hardlinked together are deleted from your local file system. When one file in a set of hardlinked files is migrated, all of the hardlinked files in the set become stub files. When the **dsmmigundelete** command recreates a stub file for a hardlinked file, the stub file has the same name as the file that was originally migrated. Stub files are not recreated for any other files that were previously in the hardlinked set of files.

Syntax



Parameters

-Expiring or -e

Recreates a stub file for a migrated file if a corresponding stub file does not exist on your local file system, whether the migrated file was marked for expiration or not. Or, it creates a stub file for a premigrated file if a corresponding original file does not exist on your local file system, whether the premigrated file was marked for expiration or not.

Enter the **dsmmigundelete** command *with* the **expiring** option if you ran reconciliation since the files were deleted.

If you do not use the **expiring** option, the HSM client recreates a stub file for a migrated file if a corresponding stub file does not exist on your local file system and the migrated file was not marked for expiration. Or, it creates a stub file for a premigrated file if a corresponding original file does not exist on your local file system, and the premigrated file was not marked for expiration.

Enter the **dsmmigundelete** command *without* the **expiring** option if you did not run reconciliation since the files were deleted.

filesystemspec

The name of the file system for which you want to recreate deleted stub files and create stub files for premigrated files that were deleted from your local file system. The default is all file systems for which space management is active. You can specify more than one file system name. If you specify several file system names, separate each name with one or more blank spaces.

Examples

Task	Command
Recreate stub files for migrated files that are marked for expiration which were accidentally deleted from the /home file system, and for files that are not marked for expiration. Reconciliation was run since the files were deleted.	<code>dsmmigundelete -expiring /home</code>
Recreate stub files for migrated files in the /home file system that are not marked for expiration. Reconciliation was not run since the files were deleted.	<code>dsmmigundelete /home</code>

dsmmode

Valid for AIX JFS environments only

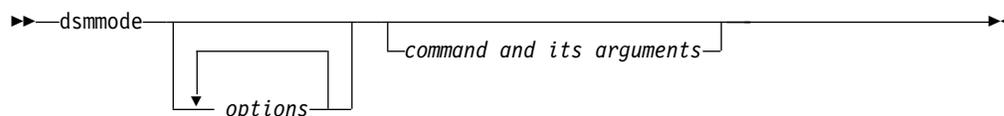
The **dsmmode** command sets one or more of the following execution modes to modify the HSM-related behavior of specific commands:

- Data access control mode. Controls access to a migrated file.
- Timestamp control mode. Sets the access time for a file to the current time when you access the file. You must have root user authority to change this execution mode.
- Out-of-space protection mode. Attempts to recover from an out-of-space condition.
- Recall mode. Recalls a migrated file when it is accessed.

The **dsmmode** command also displays the current execution modes that are in effect. If you precede another command and its arguments with the **dsmmode** command, only the execution mode or modes that you specify remain in effect for that command and any child processes of that command. Specify valid options with the **dsmmode** command for any command that follows to run. If you do not specify options, any command and its associated arguments that follow are ignored.

If you enter the **dsmmode** command without another command, it starts a new shell and the execution mode or modes that you specify remain in effect until you exit the shell. The shell that starts is the one that you specified as your default shell in `/etc/passwd`. You can nest **dsmmode** commands within the new shell.

Syntax



Parameters

options

Use any combination of the following options. If you do not specify any options, the current execution modes in effect display.

-Dataaccess=value or -dvalue

Determines how a migrated file appears to the next command that you entered. The values that you can specify are:

Value	Description
Normal <i>or</i> n	A migrated file appears to be resident, and the command can access the file. This is the default.
Zero-length <i>or</i> z	A migrated file appears to be a zero-length file.
Error <i>or</i> e	Any attempt to access a migrated file returns an EIO error.

-Timestamp=value or -tvalue

Determines whether the access time (atime) for a file is set to the current time when you access the file. Only a root user can change this mode. The values that you can specify are:

Mode	Description
Normal <i>or</i> n	The access time is set to the current time when a file is accessed. This is the default.
Preserve <i>or</i> p	The access time and inode change time (ctime) do not change when the file is read or its attributes are changed. You can specify this value to prevent access times from changing when backup programs other than the Tivoli Storage Manager backup-archive client program back up your files. Note: If you enter a command, such as the shell command <i>cat</i> , the file you are accessing is then mapped into memory. When a file is mapped, the access time will change even if you specify the preserve option.

-Outofspace=value or -ovalue

Determines how the HSM client responds to out-of-space errors that the kernel returns. The values that you can specify are:

Value	Description
Normal <i>or</i> n	The HSM client performs demand migration to avoid returning errors to the user if possible, and attempting to recover from out-of-space conditions. This is the default.
Error <i>or</i> e	The HSM client does not intercept out-of-space errors.

-Recall=value or -rvalue

Determines how the HSM client recalls a migrated file when it is accessed. You can set a recall mode for a process to normal or migrate-on-close. The following list defines the order of precedence for recall modes if the recall mode that you set for a process is different from the recall mode that you set for a migrated file accessing the process:

1. Migrate-on-close
2. Normal

The values that you can specify are:

Value	Description
Normal <i>or</i> n	When a migrated file is accessed, it is copied to its originating file system. If the file is not modified, it becomes a premigrated file. If the file is modified, it becomes a resident file. Normal is the default recall mode for a process.
Migonclose <i>or</i> m	When a migrated file is accessed, the HSM client temporarily copies it to your local file system. If the file is not modified, the file is returned to a migrated state by replacing the file with a stub file on your local file system when it is closed. If the file is modified, it remains on your local file system as a resident file.

Use the **dsmatrr** command to set a recall mode for a migrated file to normal, migrate-on-close, or read-without-recall. See Table 13 on page 75 for the recall modes that are used.

Attention: More than one process can access a file at a time. Another process can cause a file that you are accessing in migrate-on-close or read-without-recall mode to remain on your local file system as a resident or premigrated file. For example, if the recall mode that you set for a file is normal, and you access it with a process using the migrate-on-close recall mode and you do not modify the file, you would expect the file to be in a migrated state when you close it. But if someone else accesses the file at the same time with a process using normal recall mode or with a process that modifies the file in some way, the file remains on your local file system as either a resident or premigrated file, depending on the actions that the other process takes.

command

A command and any of its valid arguments for which you want the specified execution mode to remain in effect. If you do not specify a command, the HSM client starts a new shell. The shell that is started is the one that you specified as your default shell in `/etc/passwd`. The execution mode that you specify remains in effect until you exit the shell.

Examples

Task	Command
Display the current execution modes.	<code>dsmmode</code>
Change the data access mode for a grep command to ensure that only resident and premigrated files are searched, and migrated files are not accessed and recalled.	<code>dsmmode -Dataaccess=z grep \ spaceman *</code>

dsmmonitor

dsmmonitor

You must have root user authority to use this command.

The **dsmmonitor** command starts the HSM space monitor daemon if it has stopped. If you enter this command and the space monitor daemon is running, action is not taken.

Syntax

▶—dsmmonitor—▶

dsmq

You must have root user authority to use this command.

The **dsmq** command displays the following information about each recall process that is queued for processing:

- The recall ID
- The recall daemon process ID
- The start time for the recall process
- The inode number for the recalled file
- The name of the file system where the file is being recalled
- The original name of the file at the time it was migrated.

If you set the **maxrecalldaemons** option in your dsm.sys file lower than the current number of requested recalls, some recall requests will not appear in the output for this command until recall daemons are available to perform the requests. To remove a recall process from the queue, use the **dsmrm** command.

Note: If the Tivoli Storage Manager server is busy, or the connection between the HSM client and the Tivoli Storage Manager server is slow, the original name of the file might display as UNKNOWN in the output for this command. Enter **dsmq** again to view the file name.

If the recall daemon process ID (DPID) is zero, the recall is complete. You cannot remove the recall process from the queue.

Syntax

▶▶—dsmq—▶▶

Examples

Task	Command
Display the status of recall processes.	dsmq

Task**Command**

Recall all migrated files in a directory named /mfs4/user1 and all migrated files in its subdirectories. `dsmrecall -Recursive /mfs4/user1/*`

dsmrecalld

You must have root user authority to use this command.

The **dsmrecalld** command starts a recall daemon if it is not running.

If the file resides in a file system whose server stanza is different from the last processed file, a new session is started for each file that is recalled. This can happen as a result of links from one file system to another.

Note: If you enter this command while a recall daemon is running, no action is taken.

Syntax

▶▶—dsmrecalld—————▶▶

dsmreconcile

orphans are located, their names are recorded in the `.SpaceMan/orphan.stubs` file. In this instance, the **dsmautomig** command cannot run concurrently with **dsmreconcile**.

If you want to check for orphans but you do not want to perform full file system tree traversals, specify *no* on the **checkfororphans** option in your `dsm.sys` file. You can then run **dsmautomig** concurrently with **dsmreconcile**.

filesystemspec

The name of the file system to reconcile. If you do not specify a file system name, the HSM client reconciles all file systems on your workstation for which space management is active. If you enter more than one file system name, separate each name with at least one blank space.

Examples

Task	Command
Start reconciliation for all file systems for which space management is active.	<code>dsmreconcile</code>
Start reconciliation for the <code>/migfs1</code> file system.	<code>dsmreconcile /migfs1</code>
Start reconciliation for file systems <code>/home</code> and <code>/test1</code> .	<code>dsmreconcile /home /test1</code>

dsmrm

You must have root user authority to use this command.

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.

Note: After a recall process starts, enter the **dsmrm** command to stop the process. Do not use **Ctrl C** to stop a recall process.

Syntax



Parameters

recallid

The recall process ID that you want to remove from the queue.

Examples

Task	Command
Remove recall ID 10 from the queue.	dsmrm 10
Remove recall ID 5 and recall ID 6 from the queue.	dsmrm 5 6

dsmrootd

You must have root user authority to use this command.

The **dsmrootd** command starts a root daemon if it is not running. The root daemon provides non-root user support for HSM.

The following commands can be run as non-root:

- **dsmdf**
- **dsmdu**
- **dsmls**
- **dsmmigfs query**
- **dsmmighelp**
- **dsmmigrate**
- **dsmrecall**

Note: If you enter this command while a root daemon is running, no action is taken.

Syntax

▶▶—dsmrootd—————▶▶

dsmscoutd

You must have root user authority to use this command.

The **dsmscoutd** command starts the scout daemon if it has stopped. If you enter this command while a scout daemon is running, no action is taken.

Syntax

▶▶—dsmscoutd—————▶▶

dsmsetpw

You must have root user authority to use this command.

The **dsmsetpw** command changes the Tivoli Storage Manager password for your client node. To keep your password secure, enter the **dsmsetpw** command *without* your old password and new password. The system prompts you for each one. When you are prompted to enter your old and new passwords, you eliminate the possibility that another user can display your password.

If you did not register your client node with the Tivoli Storage Manager server that you contact for services, and open registration is in effect, the Tivoli Storage Manager client prompts you for registration information.

Syntax

```
▶▶ dsmsetpw [ oldpassword newpassword ] ▶▶
```

Parameters

oldpassword

The current Tivoli Storage Manager password for your client node.

newpassword

The new Tivoli Storage Manager password to set for your client node.

Examples

Task	Command
Change your current Tivoli Storage Manager password from osecret to nsecret .	<code>dsmsetpw osecretnsecret</code>

dsmwatchd

Valid on AIX GPFS and Linux86 GPFS file systems only. You must have root user authority to use this command.

Note: Do not use this command from the console. During installation, it is added to `/etc/inittab`. The **dsmwatchd** command requires a living PSSP cluster, AIX cluster Group Services, or IBM Cluster 1350 cluster and read/write access to shared HSM files within the SDR on a PSSP or IBM Cluster 1350 cluster system.

The **dsmwatchd** command manages failover activities for your local node. If failover is active on your local node, it checks the status of the **dsmrecalld** command, the **dsmmonitord** command, and the **dsmscoutd** command. If any of these daemons end, or becomes corrupted, **dsmwatchd** automatically recovers the failed daemon.

If two or more nodes within a GPFS node set participate actively in a failover environment, the **dsmwatchd** command either will take over the filesystems of a failed HSM node actively (remote is the same as within the local nodeset), or start the failover if the HSM client no longer can perform its operations locally. A node crash also can start failover. Unmounting a managed filesystem will not result in a failover. The failover environment is active by default. Use **dsmmigfs enableFailover** or **dsmmigfs disableFailover** to change the status.

In an active environment, the `dsm.opt` and `dsm.sys` files are shared between all nodes to guarantee that all nodes access the same Tivoli Storage Manager server. Even with a disabled failover environment, there is some communication between the different **dsmwatchd** commands to ensure that a GPFS filesystem is managed at least once. The `dsm.opt` and `dsm.sys` files can be different in this instance. Nodes with a different configuration no longer can participate in the failover environment. Starting failover will result in a configuration synchronization. The local `dsm.opt` and `dsm.sys` files are replaced with the corresponding SDR versions that are valid for failover. The **dsmwatchd** writes error messages to the `/dsmerror.log`. If you want **dsmwatchd** to use another `dsmerror.log` file, set the environment variable, `DSM_LOG` in `/etc/environment` accordingly. For example:

```
DSM_LOG=/usr/tivoli/tsm/client/hsm/bin/dsmerror.log
```

Attention: The synchronization operations within the failover group creates backup copies with a timestamp of the configuration files. These copies are stored in the backup directory located in the `DSM_DIR` path and in `/etc/adsm/SpaceMan/config`.

Syntax

```
▶—dsmwatchd— import—▶
```

Parameters

import

Automatically starts the following daemons on the takeover node if these daemons are not currently running: **dsmrecalld**, **dsmmonitord**, **dsmscoutd**.

dsmwatchd

Appendix A. Using the Installfsm Program

Valid on AIX JFS file systems only. You must have root user authority to run the installfsm program.

installfsm

The installfsm program installs the kernel extension and virtual file system. A vnode cache is used to improve performance on HSM-managed file systems. The larger the cache, the less often Tivoli Storage Manager must read stub file information for migrated files. If the cache is too large, it can use up all journaled file system (JFS) inode slots. This causes applications, including non-HSM-managed file systems, to receive an ENFILE return code on file system calls. The JFS inode table size varies by the amount of real memory that is installed on the machine. If you are experiencing ENFILE return codes, reduce the cache size parameter.

Start the installfsm program from the rc.adsmhsm file.

Syntax

```
▶▶—installfsm [ -l kernel_extension_name | -q kernel_extension_name | -u kernel_extension_name ] ▶▶
```

Parameters

- l Loads the kernel extension.
- q Queries to determine if the kernel extension is loaded.
- u Unloads the kernel extension.

kernel_extension_name

Specifies the kernel extension name. Enter a fully-qualified path name such as:
/usr/tivoli/tsm/client/hsm/bin/kext

Examples

Task	Command
Load the kernel extension.	installfsm -l /usr/tivoli/tsm/client/hsm/bin/kext
Unload the kernel extension.	installfsm -u /usr/tivoli/tsm/client/hsm/bin/kext

Appendix B. The .SpaceMan Directory

Tivoli Storage Manager creates file system-specific control files that reside in a hidden directory named `.SpaceMan` on each file system to which you add space management. These files also reside in the `/etc/adsm/SpaceMan` directory that Tivoli Storage Manager creates when you install the HSM client. The HSM client automatically excludes these files from space management. User action is not required to ensure that the files remain on your local file systems. Table 20 describes the file contents of this directory.

Table 20. HSM Control Files

File Name	Description
<code>config/dsmmigfstab</code>	File containing space-managed file system names and the space management settings that you selected for each file system.
<code>config/dsmmigfstab.pid</code>	A lock file for the <code>dsmmigfstab</code> file.
<code>config/dmiFSGlobalState</code>	Valid for AIX GPFS, Linux86 GPFS, Solaris VxFS, and HP-UX VxFS file systems <i>only</i> . Global and HSM-managed file system information. Related to the DMAPI interface.
<code>config/dmiFSGlobalState.pid</code>	A lock file for the <code>dmiFSGlobalState</code> file.
<code>ActiveRecallTab</code>	Active recall table.

Space Management Objects

When you add space management to your file systems, the HSM client creates control files in a hidden directory named `.SpaceMan`. Table 21 provides a brief description of objects that are stored in the `.SpaceMan` directory in each file system to which you add space management. The HSM client requires these objects for processing.

Attention: The HSM client can create other objects in the `.SpaceMan` directory during space management processing. *Do not* delete or modify the `.SpaceMan` directory, its contents, or any of the file ownerships and permissions.

Table 21. Control Files Stored in the `.SpaceMan` Directory

Files	Description
<code>orphan.stubs</code>	Records the names of any files for which a stub file exists on your local file system, but a corresponding migrated file does not exist in Tivoli Storage Manager storage.
<code>status</code>	Records space management-related statistics for your file system.
<code>premigldb.dir</code> <code>premigldb.pag</code>	Valid for AIX JFS file systems only Stores the premigrated files database that contains information about each premigrated file.
<code>dmiFSState</code>	Valid for AIX GPFS, Linux86 GPFS, Solaris VxFS, and HP-UX VxFS file systems only Stores information about the file system.

Table 21. Control Files Stored in the .SpaceMan Directory (continued)

Files	Description
logdir/	Records information during file migration or recall. It uses this information to complete any interrupted transactions, such as a system failure.
SDR/	<p data-bbox="670 363 1109 384">Valid for AIX (IBM Regatta pSeries 690)</p> <p data-bbox="670 415 1401 489">Stores information for failover that must be shared between cluster nodes that are configured as failover nodes. The directory exists in only one of the .SpaceMan directories of HSM-managed file systems.</p>
reserved/	<p data-bbox="670 510 1417 573">Valid for AIX GPFS, Linux86 GPFS, Solaris VxFS, and HP-UX VxFS file systems only</p> <p data-bbox="670 594 1157 615">Contains reserved files for demand migration.</p>

Appendix C. 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 give 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 (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

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

IBM Corporation
2Z4A/101
11400 Burnet Road
Austin, TX 78758 U.S.A.

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

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

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.

Trademarks

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both:

AIX	Operating System/400
IBM	OS/390
MVS	OS/400
OpenEdition	RISC System/6000
	System /390
	Tivoli

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc., in the United States, other countries, or both.

Microsoft, Windows, Windows NT, Windows 2000, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, and service names might be trademarks or service marks of others.

Glossary

Term definitions in this glossary pertain to the Tivoli Storage Manager library. If you do not find a term you are looking for, you can refer to the IBM Software Glossary on the Web at this address:

<http://www.ibm.com/ibm/terminology/>

This glossary might include terms and definitions from:

- The *American National Standard Dictionary for Information Systems*, ANSI X3.172-1990, copyright (ANSI). Copies can be purchased from the American National Standards Institute, 11 West 42nd Street, New York, New York 10036.
- The *Information Technology Vocabulary*, developed by Subcommittee 1, Joint Technical Committee 1, of the International Organization for Standardization and the International Electrotechnical Commission (ISO/IEC JTC2/SC1).

A

absolute. A copy group mode value that indicates a file is considered for incremental backup even though the file has not changed since the last time it was backed up. See *mode*. Contrast with *modified*.

active file system. A file system for which space management has been added. HSM can perform all space management tasks for an active file system, including automatic migration, recall, and reconciliation and selective migration and recall. Contrast with *inactive file system*.

active policy set. The policy set within a policy domain that contains the most recently activated policy. This policy set is used by all client nodes assigned to the current policy domain. See *policy set*.

active version. The most recent backup copy of a file stored in Tivoli Storage Manager storage for a file that currently exists on a file server or workstation. An active version remains active and exempt from deletion until:

- Replaced by a new backup version.
- Tivoli Storage Manager detects, during an incremental backup, that the user has deleted the original file from a file server or workstation.

administrative client. A program that runs on a file server, workstation, or mainframe. This program lets administrators monitor and control Tivoli Storage Manager servers using Tivoli Storage Manager administrator commands. Contrast with *backup-archive client*.

administrator. A user who is registered to the server as an administrator. Administrators may possess one or more privilege classes. Administrators can use the administrative client to enter Tivoli Storage Manager server commands and queries according to their privileges.

age factor. A value that determines the weight given to the age of a file when HSM prioritizes eligible files for migration. The age of the file in this case is the number of days since the file was accessed. The age factor and size factor determine migration priority for a file. See also *size factor*.

aggregate data transfer rate. Dividing the total number of bytes transferred by the elapsed processing time calculates the data transfer rate.

archive. A function permitting users to copy one or more files to a long-term storage device. Archive copies can:

- Accompany descriptive information
- Imply data compression software usage
- Be retrieved by archive date, file name, or description

Contrast with *retrieve*.

archive copy. A file or group of files residing in an archive storage pool in Tivoli Storage Manager storage.

archive copy group. A policy object containing attributes that control the generation, destination, and expiration of archived files. The archive copy group belongs to a management class.

archive retention grace period. The number of days Tivoli Storage Manager retains an archived copy when the server is unable to rebind the file to an appropriate management class.

authentication. The process of checking and authorizing a user's password before permitting user access to the Tivoli Storage Manager server. An administrator with system privilege can enable or disable authentication.

authorization rule. A specification permitting another user to either restore or retrieve a user's files from Tivoli Storage Manager storage.

automatic migration. The process HSM uses to automatically move files from a local file system to Tivoli Storage Manager storage based on options and settings chosen by a root user on your workstation. This process is controlled by the **dsmonitor** space monitor daemon. See also *threshold migration* and *demand migration*.

automatic reconciliation. The process HSM uses to reconcile your file systems at regular intervals set by a root user on your workstation. This process is controlled by the **dsmonitor** space monitor daemon. See also *reconciliation*.

automounted file system (AutoFS). A file system managed by an automounter daemon. The automounter daemon monitors a specified directory path and automatically mounts the file system to access data.

B

backup. A function permitting users to copy one or more files to a storage pool to protect against data loss. Contrast with *restore*.

backup-archive client. A program that runs on a file server, PC, or workstation and provides a means for Tivoli Storage Manager users to back up, archive, restore, and retrieve files. Contrast with *administrative client*.

backup copy group. A policy object that contains attributes controlling the generation, destination, and expiration of backup files. The backup copy group belongs to a management class.

backup retention grace period. The number of days Tivoli Storage Manager retains a backup version when the server is unable to rebind the file to an appropriate management class.

backup set. A collection of active files in your file spaces that reside on the Tivoli Storage Manager server, created by your Tivoli Storage Manager administrator and copied onto any portable media device that is supported by both the Tivoli Storage Manager server and client.

backup version. A backed up file, directory, or file space that resides in a backup storage pool in Tivoli Storage Manager storage. The *active version* is the most recent backup version. See *active version* and *inactive version*.

binding. The process of associating a file with a management class name.

C

central scheduling. A function permitting an administrator to schedule backup and archive operations from a central location. Operations can be scheduled on a periodic basis or on an explicit date.

client. A program running on a file server, PC, workstation, or terminal that requests services of another program called the server. There are two types of Tivoli Storage Manager clients: *administrative* and *backup-archive*. See *administrative client* and *backup-archive client*.

client domain. The set of drives, file systems, or volumes selected by a user for processing during a backup or archive operation.

client node. A file server or workstation registered with the server on which the backup-archive client program is installed.

client polling. A client and server communication technique where the client node queries the server for scheduled work.

client/server. A communications network architecture in which one or more programs (clients) request computing or data services from another program (the server).

client system options file. An editable file that contains communication, authorization, central scheduling, backup, archive, and space management options. The options in a client system options file are set by a root user on your workstation. The file name is **dsm.sys** and is stored in your Tivoli Storage Manager installation directory.

client user options file. A user-editable file that contains options that identify the Tivoli Storage Manager server to contact, specify backup, archive, restore, retrieve, and space management options, and set date, time, and number formats. The file name is **dsm.opt** and is stored in your Tivoli Storage Manager installation directory.

closed registration. A registration process in which an Tivoli Storage Manager administrator must register workstations as client nodes with the server. Contrast with *open registration*.

command line interface. A type of user interface where commands are specified on the command line. Contrast with *graphical user interface*.

communication method. The method by which a client and server exchange information. For Tivoli Storage Manager backup-archive clients, the method can be SNA LU6.2, or TCP/IP. See *Systems Network Architecture Logical Unit 6.2*, and *Transmission Control Protocol/Internet Protocol*.

communication protocol. A set of defined interfaces that permits computers to communicate with each other.

copy group. A policy object that contains attributes that control backup and archive file:

- Generation
- Destination
- Expiration.

Backup and archive copy groups belong to management classes. See *frequency*, *destination*, *mode*, *retention*, *serialization*, and *version*.

CPIC. Common Programming Interface Communications.

D

data access control mode. One of four execution modes provided by the **dsmmode** command. Execution modes allow you to change the space management related behavior of commands that run under **dsmmode**. The data access control mode controls whether a command can access a migrated file, sees a migrated file as zero-length, or receives an input/output error if it attempts to access a migrated file. See also *execution mode*.

data storage management application program interface (DSMAPI). A set of functions and semantics that can monitor events on files and manage and maintain the data in a file. In a Tivoli Storage Manager environment, a DSMAPI uses events to notify data management applications about operations on files, stores arbitrary attribute information with a file, supports managed regions in a file, and uses DSMAPI access rights to control access to a file object.

default management class. A management class assigned to a policy set. This class is used to govern backed up or archived files when a user does not explicitly associate a file with a specific management class through the include-exclude list.

demand migration. The process HSM uses to respond to an out-of-space condition on a file system. HSM migrates files to Tivoli Storage Manager storage until space usage drops to the low threshold set for the file system. If the high threshold and low threshold are the same, HSM attempts to migrate one file.

destination. A copy group attribute that specifies the storage pool in which to back up or archive a file. At installation, Tivoli Storage Manager provides two storage destinations named **backuppools** and **archivepool**.

domain. See *policy domain* or *client domain*.

drag. Move the mouse while holding down the mouse button, thus moving the selected object.

drag-and-drop. Move (drag) an object on top of another object and release the mouse button, thus relocating the object.

dsm.opt file. See *options file*. See also *client user options file*. Also called client options file.

dsm.sys file. See *options file* or *client system options file*.

dynamic. A copy group serialization value that specifies Tivoli Storage Manager accept the first attempt to back up or archive an object, regardless of any changes made during backup or archive processing. See *serialization*. Contrast with *shared dynamic*, *shared static*, and *static*.

E

error log. A text file written on disk that contains Tivoli Storage Manager processing error messages. These errors are detected and saved by the Tivoli Storage Manager server.

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. You can exclude a file from backup and space management, backup only, or space management only.

execution mode. A mode that controls the space management related behavior of commands that run under the **dsmmode** command. The **dsmmode** command provides four execution modes — a data access control mode that controls whether a migrated file can be accessed, a time stamp control mode that controls whether the access time for a file is set to the current time when the file is accessed, an out-of-space protection mode that controls whether HSM intercepts an out-of-space condition on a file system, and a recall mode that controls whether a file is stored on your local file system when accessed, or stored on your local file system only while it is being accessed, and then migrated back to Tivoli Storage Manager storage when it is closed.

expiration. The process in which files are identified for deletion because their expiration date or retention period has passed. Backed up or archived files are marked for deletion based on the criteria defined in the backup or archive copy group.

expiring file. A migrated or premigrated file that has been marked for expiration and removal from Tivoli Storage Manager 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. It expires and is removed from Tivoli Storage Manager storage after the number of days specified with the **migfileexpiration** option have elapsed.

F

file access time. The user access time (atime) for file migration eligibility and priority. HSM root users can preserve the atime (and ctime) of files on an HSM active file system using a timestamp preserve dsmmode shell.

file age. For migration prioritization purposes, the number of days since a file was last accessed.

file server. A dedicated computer and its peripheral storage devices connected to a local area network that stores both programs and files shared by users on the network.

file size. For migration prioritization purposes, the size of a file in 1-KB blocks.

file space. A logical space on the Tivoli Storage Manager server that contains a group of files. In Tivoli Storage Manager, users can restore, retrieve, or delete file spaces from Tivoli Storage Manager storage. A file space for systems:

- **Windows**— Logical partition identified by a volume label.
- **UNIX** — Logical space that contains a group of files backed up or archived from the same file system, or part of a file system defined with the virtualmountpoint option in the client system options file.

file state. The state of a file that resides in a file system to which space management has been added. A file can be in one of three states — resident, premigrated, or migrated. See also *resident file*, *premigrated file*, and *migrated file*.

file system migrator (FSM). A kernel extension that is mounted over an operating system file system when space management is added to the file system. The file system migrator intercepts all file system operations and provides any space management support that is required. If no space management support is required, the operation is performed by the operating system file system.

file system state. The state of a file system that resides on a workstation on which Tivoli Storage Manager HSM is installed. A file system can be in one of these states— native, active, inactive, or global inactive.

FSM. File system migrator.

frequency. A copy group attribute that specifies the minimum interval, in days, between incremental backups.

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. An archive copy of a file that might not accurately reflect what is currently in the file because Tivoli Storage Manager archived the file while the file was being modified.

G

generate password. Processing that stores a new password in an encrypted password file when the old password expires. Automatic generation of a password prevents password prompting. Password generation can be set in the options file (passwordaccess option). See *options file*.

gigabyte (GB). (1) One billion (10^9) bytes. (2) When referring to memory capacity, 1 073 741 824 in decimal notation.

global inactive state. The state of all file systems to which space management has been added when space management is globally deactivated for a client node. When space management is globally deactivated, HSM cannot perform migration, recall, or reconciliation. However, a root user can update space management settings and add space management to additional file systems. Users can access resident and premigrated files.

GPFS node set. A set of AIX SP or Linux86 nodes that can mount a defined group of GPFS file systems.

graphical user interface (GUI). A user interface consisting of a visual metaphor representing a real-world scene, often looks like a desktop. A graphical user interface includes:

- A combination of graphics and icons
- The object-action paradigm
- Use of pointing devices, menu bars, and overlapping windows

Contrast with *command line interface*. See *windowed interface*.

GUI. Graphical user interface.

H

hierarchical storage management client. A program that runs on a workstation or file server to provide space management services. The hierarchical storage management client automatically migrates eligible files to Tivoli Storage Manager storage to maintain specific levels of free space on local file systems. Automatic recalls are made for migrated files when they are accessed. Users are also permitted to migrate and recall specific files.

high threshold. The percentage of space usage on a local file system at which HSM automatically begins migrating eligible files to Tivoli Storage Manager storage. A root user sets this percentage when adding space management to a file system or updating space management settings. Contrast with *low threshold*.

HSM. Hierarchical Storage Management.

I

image. A full file system or raw logical volume backup as a single object.

inactive file system. A file system for which you have deactivated space management. When space management is deactivated for a file system, HSM cannot perform migration, recall, or reconciliation for the file system. However, a root user can update space management settings for the file system, and users can access resident and premigrated files. Contrast with *active file system*.

inactive version. A copy of a backup file in Tivoli Storage Manager storage that either is not the most recent version, or the corresponding original file was deleted from the client file system. Inactive backup versions are eligible for expiration according to the management class assigned to the file.

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 *include-exclude list*.

include-exclude list. A list of include and exclude 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. The include-exclude list is defined in one or more include-exclude files or in the . . See *options file*.

incremental backup. A function that permits user to back up new or changed files or directories from a client domain or from specified . These are not excluded in the include-exclude list and meet the requirements for frequency, mode, and serialization as defined by a backup copy group of the management class assigned to each file. Contrast with *selective backup*.

inode. A data structure that describes the individual files in an operating system. There is one inode for each file. The number of inodes in a file system, and therefore the maximum number of files a file system can contain, is set when the file system is created. Hardlinked files share the same inode.

inode number. A number that specifies a particular inode in a file system.

L

LAN. Local area network.

LAN-free data transfer. The movement of client data between the client and a storage device over a SAN, bypassing the LAN.

leader data. Leading bytes of data from a migrated file that are stored in the file's corresponding stub file on the local file system. The amount of leader data stored in a stub file depends on the stub size specified. The required data for a stub file consumes 511 bytes of space. Any remaining space in a stub file is used to store leader data. If a process accesses only the leader data and does not modify that data, HSM does not need to recall the migrated file back to the local file system.

Local Area Network (LAN). A variable-sized communications network placed in one location. LAN connects servers, PCs, workstations, a network operating system, access methods, and communications software and links.

logical volume backup. A back up of a file system or logical volume as a single object

Loopback Virtual File System (LOFS). A file system created by mounting a directory over another local directory, also known as mount-over-mount. A LOFS can also be generated using an automounter.

low threshold. A percentage of space usage on a local file system at which HSM automatically stops migrating files to Tivoli Storage Manager storage during a threshold or demand migration process. A root user sets this percentage when adding space management to a file system or updating space management settings. Contrast with *high threshold*.

M

management class. A policy object that is a named collection of copy groups. A management class is associated with a file to specify how the server should manage backup versions or archive copies of workstation files. The space management attributes contained in a management class assigned to a file determine whether the file is eligible for automatic or selective migration. See *binding* and *copy group*.

migrate-on-close recall mode. A mode that causes HSM to recall a migrated file back to its originating file system only temporarily. If the file is not modified, HSM returns the file to a migrated state when it is closed. However, if the file is modified, it becomes a resident file. You can set the recall mode for a migrated file to migrate-on-close by using the **dsmatrr** command, or set the recall mode for a specific execution of a command or series of commands to migrate-on-close by using the **dsmmode** command. Contrast with *normal recall mode* and *read-without-recall recall mode*.

migrated file. A file that has been copied from a local file system to Tivoli Storage Manager storage and replaced with a stub file on the local file system. Contrast with *resident file* and *premigrated file*.

migration. The process of copying a file from a local file system to Tivoli Storage Manager storage and replacing the file with a stub file on the local file system. See also *threshold migration*, *demand migration*, and *selective migration*.

mode. A copy group attribute that specifies whether a backup file should be created for a file that was not modified since the last time the file was backed up. See *absolute* and *modified*.

modified. A backup copy group attribute indicating a file is considered for backup only if the file has been changed since the last backup. A file is considered changed if the date, size, owner, or permissions have changed. See *absolute* and *mode*.

N

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.

native file system. A file system to which you have not added space management. If space management is not added to a file system, the HSM client cannot provide space management services.

network data transfer rate. The data transfer rate calculated by dividing the total number of bytes transferred by the data transfer time. For example, the time spent transferring data over the network.

node. See *client node*.

node name. A unique name used to identify a workstation, file server, or PC to the server.

normal recall mode. A mode that causes HSM to copy a migrated file back to its originating file system when it is accessed. If the file is not modified, it becomes a premigrated file. If the file is modified, it becomes a resident file. Contrast with *migrate-on-close recall mode* and *read-without-recall recall mode*.

NS/DOS. Networking Services/DOS.

O

open registration. A registration process in which users can register their own workstations or PCs as client nodes with the server. Contrast with *closed registration*.

options file. A file that contains processing options.

- **dsm.opt**

Non-UNIX — Identifies Tivoli Storage Manager servers, specifies communication methods, defines scheduling options, selects backup, archive, restore, and retrieve options. Also called the client options file.

UNIX — Identifies the Tivoli Storage Manager server to contact, specifies backup, archive, restore, and retrieve options. Also called the client users options file.

- **dsm.sys**

UNIX — Contains stanzas describing Tivoli Storage Manager servers to contact for services. These stanzas also specify communication methods, backup and archive options, and select scheduling options. Also called the client system options file.

orphaned stub file. A stub file for which no migrated file can be found on the Tivoli Storage Manager server your client node is currently contacting for space management services. A stub file can become orphaned, for example, if you modify your client system options file to contact a different server for space management than the one to which the file was migrated.

out-of-space protection mode. One of four execution modes provided by the **dsmmode** command. Execution modes allow you to change the HSM-related behavior of commands that run under **dsmmode**. The out-of-space protection mode controls whether HSM intercepts out-of-space conditions. See also *execution mode*.

originating file system. The file system from which a file was migrated. When a file is recalled using normal or migrate-on-close recall mode, it is always returned to its originating file system.

owner. The owner of backup-archive files sent from a multi-user client node, such as AIX.

P

partial file recall mode. A recall mode that causes HSM to read just a portion of a migrated file from Tivoli Storage Manager storage, as requested by the application accessing the file. Partial file recall is valid for AIX GPFS HSM only and only available for transparent, not selective recalls.

pattern-matching character. See *wildcard character*.

policy domain. A policy object that contains one or more policy sets. Client nodes are associated with a policy domain. See *policy set*, *management class*, and *copy group*.

policy set. A policy object that contains a group of management class definitions that exist for a policy domain. At any one time, there can be many policy sets within a policy domain, but only one policy set can be active. See *active policy set* and *management class*.

premigrated file. A file that has been copied to Tivoli Storage Manager storage, but has not been replaced with a stub file on the local file system. An identical copy of the file resides both on the local file system and in Tivoli Storage Manager storage. When free space is needed, HSM verifies that the file has not been modified and replaces the copy on the local file system with a stub file. HSM premigrates files after automatic migration is complete if there are additional files eligible for migration, and the premigration percentage is set to allow premigration. Contrast with *migrated file* and *resident file*.

premigrated files database. A database that contains information about each file that has been premigrated to Tivoli Storage Manager storage. The database is stored in a hidden directory named **.SpaceMan** in each file system to which space management has been added. HSM updates the premigrated files database whenever it premigrates and recalls files and during reconciliation.

premigration. The process of copying files that are eligible for migration to Tivoli Storage Manager storage, but leaving the original file intact on the local file system.

premigration percentage. A space management setting that controls whether the next eligible candidates in a file system are premigrated following threshold or demand migration. The default for premigration percentage is the difference between the percentage specified for the high threshold and the percentage specified for the low threshold for a file system.

progress indicator. A control used to inform a user about the progress of a process.

Q

quota. The total number of megabytes of data that can be migrated and premigrated from a file system to Tivoli Storage Manager storage. The default for quota is the same number of megabytes as allocated for the file system itself.

R

raw logical volume. A portion of a physical volume which is comprised of unallocated blocks and has no Journaled File System (JFS) definition. A raw logical volume is read/write accessible only through low level I/O functions.

read-without-recall recall mode. A mode that causes HSM to read a migrated file from Tivoli Storage Manager storage without storing it back on the local file system. The last piece of information read from the file is stored in a buffer in memory on the local file system. However, if a process that accesses the file writes to or modifies the file or uses memory mapping, HSM copies the file back to the local file system. Or, if the migrated file is a binary executable file, and the file is executed, HSM copies the file back to the local file system. You can change the recall mode for a migrated file to read-without-recall by using the **dsmatr** command. Contrast with *normal recall mode* and *migrate-on-close recall mode*.

recall. The process of copying a migrated file from Tivoli Storage Manager storage back to its originating file system. See also *transparent recall*, *selective recall*, and *recall mode*.

recall mode. 1) One of four execution modes provided by the **dsmmode** command. Execution modes allow you to change the HSM-related behavior of commands that run under **dsmmode**. The recall mode controls whether an unmodified, recalled file is returned to a migrated state when it is closed. 2) A mode assigned to a migrated file with the **dsmatr** command that determines how the file is processed when it is recalled. It determines whether the file is stored on the local file system, is migrated back to Tivoli Storage Manager storage when it is closed, or is read from Tivoli Storage Manager storage without storing it on the local file system.

reconciliation. The process of synchronizing a file system to which you have added space management with the Tivoli Storage Manager server you contact for space management services and building a new migration candidates list for the file system. HSM performs reconciliation automatically at intervals specified with the **reconcileinterval** option in your client system options file. A root user can also start reconciliation manually at any time.

registration. The process of identifying a client node or administrator to the server by specifying a user ID, password, and contact information. For client nodes, a policy domain, compression status, and deletion privileges are also specified.

resident file. A file that resides on a local file system. It has not been migrated or premigrated, or it has been recalled from Tivoli Storage Manager storage and modified. When first created, all files are resident. Contrast with *premigrated file* and *migrated file*.

restore. A function that permits users to copy a version of a backup file from the storage pool to a workstation or file server. The backup copy in the storage pool is not affected. Contrast with *backup*.

retention. The amount of time, in days, that inactive backed up or archived files are retained in the storage pool before they are deleted. The following copy group attributes define retention: retain extra versions, retain only version, retain version.

retrieve. A function permitting users to copy an archived file from the storage pool to the workstation or file server. The archive copy in the storage pool is not affected. Contrast with *archive*.

root user (UNIX). The authority level for a root user permits this user to do TSM-authorized tasks for Tivoli Storage Manager.

S

scheduling mode. The type of scheduling operation for the client-server node. Tivoli Storage Manager supports two scheduling modes: client-polling and server-prompted.

scroll. Move through a list of items in a window by operating the scrollbars with the mouse cursor.

select. Choose an item from a list or group of items.

selective backup. A function permitting users to back up specified files. These files are not excluded in the include-exclude list and meet the requirement for serialization in the backup copy group of the management class assigned to each file. Contrast with *incremental backup*.

selective migration. The process of copying user-selected files from a local file system to Tivoli Storage Manager storage and replacing the files with stub files on the local file system. Contrast with *threshold migration* and *demand migration*.

selective recall. The process of copying user-selected files from Tivoli Storage Manager storage back to a local file system. Contrast with *transparent recall*.

serialization. A copy group attribute that specifies whether a file can be modified during a backup or archive operation. See *static*, *dynamic*, *shared static*, and *shared dynamic*.

server. A program running on a mainframe, workstation, or file server that provides shared services such as backup and archive to other various (often remote) programs (called clients).

server-prompted scheduling. A client-server communication technique where the server contacts the client node when tasks need to be done.

session. A period of time in which a user can communicate with a server to perform backup, archive, restore, or retrieve requests.

shared dynamic. A Tivoli Storage Manager copy group serialization mode. This mode specifies if a file changes during backup or archive and continues to change after a number of retries. The last retry commits the file to the Tivoli Storage Manager server whether or not the file changed during backup or archive. Contrast with *dynamic*, *shared static*, and *static*.

shared static. A copy group serialization value specifying that a file must not be modified during a backup or archive operation. Tivoli Storage Manager attempts to retry the operation a number of times. If the file is in use during each attempt, the file is not backed up or archived. See *serialization*. Contrast with *dynamic*, *shared dynamic*, and *static*.

share point. A drive or directory on Windows NT, 95, 98, or Windows Me whose files are available for shared access across a network. The share point name is part of a UNC name. See *Universal Naming Convention (UNC)* name.

shift-click. Click on an item while pressing the Shift key.

size factor. A value that determines the weight given to the size of a file when HSM prioritizes eligible files for migration. The size of the file in this case is the size in 1-KB blocks. The size factor is used with the age factor to determine migration priority for a file. See also *age factor*.

space management. The process of keeping sufficient free storage space available on a local file system for new data and making the most efficient and economical use of distributed storage resources.

space management attributes. Attributes contained in a management class that specify:

- Whether automatic or selective migration is permitted for a file.
- How many days must elapse since the file was last accessed before it is eligible for automatic migration.
- Whether a backup version of a file must exist before the file can be migrated.
- The Tivoli Storage Manager storage pool to which files are to migrate.

space management settings. Settings that specify the stub file size, quota, age factor, size factor, high threshold, low threshold, and the premigration percentage for a file system. A root user selects space management settings when adding space management to a file system or when updating space management.

space monitor daemon. A Tivoli Storage Manager HSM 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. How often the space monitor daemon checks space usage is determined by the *checkthresholds* option in your client system options file. In addition, the space monitor daemon starts reconciliation for your file systems at the intervals specified with the *reconcileinterval* option in your client system options file.

stabilized file space. A file space that exists on the server but not on the client. This situation can arise in at least two instances:

1. A drive is removed from a client workstation
2. A file space is renamed on the server

Stabilized file spaces remain on the server until deleted by the user or administrator. Files and directories can be restored and retrieved from a stabilized file space. However, it is not possible to back up or archive data to a stabilized file space.

static. A copy group serialization value specifying that a file must not be modified during a backup or archive operation. If the file is in use during the first attempt, Tivoli Storage Manager will not back up or archive the file. See *serialization*. Contrast with *dynamic*, *shared dynamic*, and *shared static*.

storage area network (SAN). A high-speed communications network optimized for storage.

storage agent. A program that enables IBM Tivoli Storage Manager to back up and restore client data directly to and from SAN-attached storage.

storage pool. A named set of storage volumes used as the destination of backup, archive, or migrated copies.

stub file. A file that replaces the original file on a local file system when the file is migrated to Tivoli Storage Manager storage. A stub file contains the information necessary to recall a migrated file from Tivoli Storage Manager storage. It also contains additional information that HSM can read to eliminate the need to recall a migrated file.

stub file size. The size of a file that replaces the original file on a local file system when the file is migrated to Tivoli Storage Manager storage. The size 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.

system drive or partition. On Windows NT, the drive or partition on which Windows NT is installed.

T

TCA. Trusted Communications Agent

TCP/IP. Transmission Control Protocol/Internet Protocol.

threshold migration. The process of moving files from a local file system to Tivoli Storage Manager storage based on the high and low thresholds defined for the file system. Threshold migration is started automatically by HSM and can be started manually by a root user. Contrast with *demand migration* and *selective migration*.

timeout. A time event involving:

- An event that happens at the end of a predetermined period of time that began at the happening of another specified event.
- A time interval allotted for certain operations to happen. For example, response to polling or addressing before system operation is interrupted and must be restarted.
- A terminal feature that logs off a user if an entry is not made within a specified period of time.

timestamp control mode. One of four execution modes provided by the **dsmmode** command. Execution modes allow you to change the space management related behavior of commands that run under **dsmmode**. The timestamp control mode controls whether commands preserve the access time for a file or set it to the current time. See also *execution mode*.

IBM Tivoli Storage Manager. A client-server licensed program product that provides storage management and data access services to customers in a multivendor computer environment.

Transmission Control Protocol/Internet Protocol (TCP/IP). A standard set of communication protocols that supports peer-to-peer connectivity of functions for both local and wide-area networks.

transparent recall. The process HSM uses to automatically recall a file back to your workstation or file server when the file is accessed. The recall mode set for a file and the recall mode set for a process that accesses the file determine whether the file is stored back on the local file system, stored back on the local file system only temporarily if it is not modified, or read from Tivoli Storage Manager storage without storing it back on the local file system. See also *recall mode*. Contrast with *selective recall*.

Trusted Communications Agent (TCA) (UNIX). A program that can handle the sign-on password protocol when password access is generated. The main process (for example, **dsm**, **dsmc**) makes a run time decision based on the password access option setting, the user ID, and the executables' access privileges to run this program. The file that contains this program must have the 's' bit set in its mode field and the owner must be root.

TSM authorized user. A user who has administrative authority for the Tivoli Storage Manager client on a workstation. This user changes passwords, performs open registrations, and deletes file spaces.

U

Universal Naming Convention (UNC) name. A name used on Windows to access a drive or directory containing files shared across a network. The UNC name includes the machine name and a share point name that represents the shared drive or directory. See *share point*.

V

version. Storage management policy may allow back-level copies of backed up objects to be kept at the server whenever an object is newly backed up. The most recent backed up copy is called the "active" version. Earlier copies are "inactive" versions. The following backup copy group attributes define version criteria: versions data exists, and versions data deleted.

W

wildcard character. An asterisk (*) or question mark (?) character used to represent multiple (*) or single (?) characters when searching for various combinations of characters in alphanumeric and symbolic names.

windowed interface. A type of user interface that is either a graphical user interface or a text-based interface. The text-based interface maintains a close affinity to the graphical user interface, including action bars and their associated pull-down menus and windows. See *graphical user interface*.

workstation. A programmable high-level workstation (usually on a network) with its own processing hardware such as a high-performance personal computer. In a local area network, a personal computer that acts as a single user or client. A workstation can also be used as a server.

Index

Special characters

- .SpaceMan directory
 - contents of 161
 - creating 44
 - location of 161

A

- activating space management
 - for a file system 127
- adding space management to a file system 44, 124
- age factor
 - description of 47
 - setting 47, 124
- AIX GPFS
 - client components 6
 - communication methods 9
 - environment requirements 8, 16
 - hardware requirements 8
 - installing 6
 - operating system requirements 8
 - software requirements 8
- AIX JFS
 - client components 11
 - communication methods 11
 - environment requirements 11
 - installfs 159
 - installing 11
 - operating system requirements 11
 - software requirements 11
- archiving migrated files 69
- automatic migration
 - defined 1
 - demand 56
 - description of 56
 - file eligibility 57
 - file prioritization 47, 56, 57
 - threshold 56
 - threshold migration 57, 118
- automatic reconciliation
 - description of 79
 - frequency 79
 - starting 79
 - tasks performed 79
- automignonuse attribute 29

B

- back up files
 - after migrating 62
 - before migrating 61
- backing up
 - migrated files 61
 - premigrated files 61
- backup and migration
 - server is different 66
 - server is the same 65
- Bourne and Korn shell variables,
 - setting 35

- building a migration candidates list 57, 80, 151

C

- C shell variables, setting 35
- cancelling a recall process 153
- candidates file 161
- candidatesinterval option 28, 57, 92
- central scheduling
 - description of 87
 - setting up 87
 - starting the client scheduler 87
- changing
 - execution modes for commands 143
 - recall mode for a command 143
 - recall mode for a migrated file 116
- changing your password 22
- checkfororphans option 28, 93
- checkthresholds option 28, 94
- client and server options, displaying 137
- client components
 - for AIX GPFS 6
 - for AIX JFS 11
 - for HP-UX 14
 - for Linux86 GPFS 16
 - for Solaris 19
- client node
 - registering with server 21
- client options
 - optionformat 108
 - overlaprecall 109
 - restoremigstate 111
- client scheduler, starting 87
- client system options
 - candidatesinterval 28, 92
 - checkfororphans 93
 - checkthresholds 28, 94
 - compression 28
 - defaultserver 28
 - errorprog 28, 97
 - kernelmessages 28, 98
 - maxcandprocs 99
 - maxmigrators 28, 100
 - maxrecalldaemons 28, 101
 - maxreconcileproc 28, 102
 - maxthresholdproc 28, 103
 - migfileexpiration 28, 104
 - migrateserver 28, 105
 - minmigfilesize 106
 - minrecalldaemons 28, 107
 - reconcileinterval 28, 110
- client system options file
 - setting space management options 26
- client user options file
 - copying or modifying 29
- client-server options
 - displaying 25
- closed registration, using 21

- commands
 - descriptions of 114
 - dsmattr 116
 - dsmautomig 118
 - dsmdf 120
 - dsmdu 121
 - dsmls 122
 - dsmmigfs 124, 127, 128, 130, 131, 133, 134, 135
 - dsmmighelp 136
 - dsmmigquery 137
 - dsmmigrate 139
 - dsmmigundelete 63, 141
 - dsmmode 143
 - dsmmonitor 83, 146
 - dsmq 147
 - dsmrecall 148
 - dsmrealld 83, 150
 - dsmreconcile 151
 - dsmrm 153
 - dsmrootd 154
 - dsmscoutd 155
 - dsmsetpw 156
 - entering 113
 - using 113
- communication methods
 - for AIX GPFS 9
 - for AIX JFS 11
 - for Linux86 GPFS 17
 - HP-UX 14
 - Solaris 19
- compression option 28, 95
- compression processing 95
- creating stub files for premigrated files 141
- customer support
 - contacting xi

D

- daemons
 - dsmmonitor 146
 - dsmrealld 150
 - dsmrootd 154
 - dsmscoutd 57, 155
 - dsmwatchd 157
 - space management 82
 - stopping 84
- data access control mode 143
- deactivating space management 127
 - from a file system 51
 - globally 51, 133
- defaultserver option 28, 96
- deleted stub files, restoring 63, 141
- demand migration 56
 - defined 1
- disable failover 130
- disaster recovery 67
- disk loss, restoring 67
- disk space requirements
 - for AIX GPFS 8

- disk space requirements (*continued*)
 - for AIX JFS 11
 - for Linux86 GPFS 16
 - HP-UX 14
 - Solaris 19
- display option information 25
- displaying
 - client and server options 137
 - execution modes 143
 - file system capacity 85
 - file system state 85, 120
 - file system type 85
 - list of files 122, 137
 - local free space 85, 120
 - local premigrated space 85, 120
 - migration candidates list 137
 - number of migrated files 85, 120
 - number of premigrated files 85, 120
 - number of resident inodes 85, 120
 - ordered recall list 137
 - recall mode for a migrated file 116
 - server migrated space 85, 120
 - space management information 85, 120, 121, 122, 131, 137
 - space management settings 85, 131
 - space usage information 85, 120, 121
- displaying online help 114
- DSM_CONFIG environment variable, setting 34
- DSM_DIR environment variable, setting 34
- DSM_LOG environment variable, setting 34
- dsm.opt file
 - copying 29
 - modifying 29
- dsm.sys file
 - setting space management options 26
- dsmattr command 77, 116
- dsmautomig command 118
 - start threshold migration 57
- dsmdf command 120
- dsmdu command 121
- dsmls command 122
- dsmmigfs command 134, 135
 - adding space management 124
 - deactivating space management 127
 - displaying settings 131
 - exporting 128
 - globally deactivating space management 133
 - globally reactivating space management 133
 - importing 128
 - manage recovery 130
 - reactivating space management 127
 - reset in the SDR 130
 - update files in SDR 130
 - updating space management settings 124
- dsmmigfstab file
 - contents of 44
 - location of 44
- dsmmighelp command 136
- dsmmigquery command 82, 137
- dsmmigrate command 139

- dsmmigrate command (*continued*)
 - selective migration 60
- dsmmigundelete command 63, 141
- dsmmode command 143
 - using the 77
- dsmmonitord command 83, 146
- dsmq command 147
- dsmrecall command 77, 148
- dsmrecalld command 83, 150
- dsmreconcile command 82, 151
- dsmrm command 153
- dsmrootd command 154
- dsmscoutd command 155
- dsmscoutd daemon 57
- dsmsetpw command 156
- dsmwatchd command 157

E

- enable failover 130
- entering commands 113
- entering options with commands 113
- environment
 - HACMP 36
 - IBM Cluster 1350 39
 - limitations of, HACMP 37
 - limitations of, IBM Cluster 1350 40
- environment requirements
 - for AIX GPFS 8, 16
 - for AIX JFS 11
 - HP-UX 14
 - Solaris 19
- environment variables
 - DSM_CONFIG 34
 - DSM_DIR 34
 - DSM_LOG 34
- errorprog option 28, 97
- exclude.backup option 32
- exclude.file.spacemgmt option 32
- exclude.spacemgmt option 32
- excluding files
 - from backup 32
 - from migration 32
- execution modes
 - data access control mode 143
 - displaying 143
 - out-of-space protection mode 143
 - recall mode 143
 - setting 143
 - timestamp control mode 143
- exporting file systems
 - using NFS servers 50
- exporting migrated files 51

F

- faster migration 2
- file systems
 - adding space management 48
 - adding space management to nested 49
 - adding space management to parent 49
 - mounting automatically 43
 - mounting manually 43
 - reconciling 79

- file systems (*continued*)
 - remounting 43
- files
 - backing up 61
 - for AIX GPFS 6
 - for AIX JFS 11
 - for HP-UX 14
 - for Linux86 GPFS 16
 - for Solaris 19
 - migrating 61
 - migration 1, 55
 - moving to another directory 6
 - restoring 61
- files systems
 - limitations when exporting with NFS servers 50
- fsm
 - mounting automatically 43
 - mounting manually 43
 - remounting 43

G

- globally deactivate space management 51
- globally deactivating space management 133
- globally reactivate space management 52
- globally reactivating space management 133
- group numbers, AIX JFS 13

H

- HACMP for AIX 36
- hardware requirements
 - for AIX GPFS 8
 - for AIX JFS 11
 - for Linux86 GPFS 16
 - HP-UX 14
 - Solaris 19
- high threshold 46
 - setting 124
- HP-UX
 - client components 14
 - communication methods 14
 - considerations for installing HSM 14
 - hardware requirements 14
 - installing 14
 - operating system requirements 14
 - software requirements 14
- HSM client
 - root user tasks 1
 - summary 1

I

- IBM Cluster 1350 for Linux 39
- importing migrated files 51
- include option 32
- include-exclude list
 - excluding files from backup 32
 - including files for backup 32
- include.file option 32

- including files
 - for backup 32
 - for migration 32
- install from the command line
 - AIX GPFS 9
- installation requirements
 - for AIX GPFS 6
 - for AIX JFS 11
 - for HP-UX 14
 - for Linux86 GPFS 16
 - for Solaris 19
 - HP-UX 14
 - Solaris 20
- installfsm program 160
 - AIX JFS 159
- installing
 - kernel extension 160
 - virtual file system 160
- installing HSM
 - HP-UX 14
 - on AIX GPFS 5, 6
 - on AIX JFS 5, 11
 - on HP-UX 14
 - on Linux86 GPFS 16
 - on Solaris 19
 - on VxFS 5
 - Solaris 19, 20

K

- kernel extension
 - installing on AIX JFS 160
- kernelmessages option 28, 98

L

- LAN-free data transfer
 - options for 35
 - overview 35
 - prerequisites for 35
- Linux86 GPFS
 - client components 16
 - communication methods 17
 - hardware requirements 16
 - installing 16
 - operating system requirements 17
 - software requirements 16
- list of files, displaying 137
- logdir directory 161
- low threshold 46
 - setting 124

M

- manage recovery 130
- management class attributes
 - automignonuse 29
 - migdestination 29
 - migrequiresbkup 29
 - spacemgtechnique 29
- management classes
 - assign to files 29
 - displaying information about 29
- maxcandidates
 - setting 124
- maxcandprocs option 57, 99

- maxmigrators option 28, 100
- maxrealldaemons option 28, 101
- maxreconcileproc option 28, 102
- maxthresholdproc option 28, 103
- memory requirements
 - for AIX GPFS 8
 - for AIX JFS 11
 - for Linux86 GPFS 16
 - HP-UX 14
 - Solaris 19
- migdestination attribute 29
- migfileexpiration 28
- migfileexpiration option 80, 104
- migrate-on-close recall
 - defined 2
- migrate-on-close recall mode
 - description of 72, 116
 - setting for a command 143
 - setting for a migrated file 116
- migrated file
 - deleting from Tivoli Storage Manager storage 79
 - setting a recall mode for 116
- migrated files
 - recalling 71
- migrateserver option 28, 105
- migrating files
 - automatically 56, 57, 118
 - manually 57
 - recording information about 161
 - selectively 58, 139
- migrating files from a local file system 1
- migrating your files 55
- migration
 - automatic 55, 56
 - demand 1
 - threshold 1
 - defined 1
 - demand 56
 - file eligibility 57, 58
 - file prioritization 47, 56, 57
 - minimum file size for 45
 - premigration 2
 - selective 1, 55
 - selective migration 58, 139
 - threshold 56
 - threshold migration 57, 118
 - types of 1
- migration candidates list
 - building 57, 80
 - description of 161
 - displaying 137
- migration server, setting 28
- migrequiresbkup attribute 29
- minimum file size for migration 45
- minmigfilesize option 57, 106
- minrealldaemons option 28, 107
- modifying your include-exclude options
 - file 32
- mounting file systems
 - AIX GPFS file system 43
 - AIX JFS file system 43
 - file system migrator (FSM) 43
 - Linux86 GPFS file system 43
 - Veritas file system (VxFS) 43
- moving files to another directory 6

N

- nested file systems
 - adding space management to 49
- NFS
 - limitations when exporting space-managed file systems 50
- NFS servers
 - exporting files using 50
- normal recall
 - defined 3
- normal recall mode
 - description of 116
 - setting for a command 143
 - setting for a migrated file 116

O

- online help
 - displaying for commands 114, 136
- open registration, using 21
- operating system requirements
 - AIX GPFS 8
 - for AIX JFS 11
 - HP-UX 14
 - Linux86 GPFS 17
 - Solaris 19
- optionformat option 108
- options
 - candidatesinterval 92
 - checkfororphans 93
 - checkthresholds 94
 - client system options 26
 - compression 95
 - defaultserver 96
 - display information 25
 - entering with commands 113
 - errorprog 97
 - exclude.backup 32
 - exclude.file.spacemgmt 32
 - exclude.spacemgmt 32
 - include 32
 - include-exclude options 32
 - include.file 32
 - kernelmessages 98
 - maxcandprocs 99
 - maxmigrators 100
 - maxrealldaemons 101
 - maxreconcileproc 102
 - maxthresholdproc 103
 - migfileexpiration 104
 - migrateserver 105
 - minmigfilesize 106
 - minrealldaemons 107
 - optionformat 108
 - overlaprecall 109
 - reconcileinterval 110
 - restoremigstate 111
 - setting space management 25
- options, client system
 - candidatesinterval 28
 - checkfororphans 28
 - checkthresholds 28
 - compression 28
 - defaultserver 28
 - errorprog 28
 - kernelmessages 28

- options, client system *(continued)*
 - maxrecalldaemons 28
 - maxreconcileproc 28
 - maxthresholdproc 28
 - migfileexpiration 28
 - migrateserver 28
 - minrecalldaemons 28
 - reconcileinterval 28
- options, displaying 25
- ordered recall list, displaying 137
- orphan.stubs file 79, 161
- orphaned stub file
 - identifying 161
 - resolving 80
- out-of-space protection mode 143
- overlaprecall option 109
- overview of HSM 1

P

- parent file system
 - adding space management to 49
- partial file recall
 - defined 3
- partial file recall mode
 - description of 74
- partial recall mode
 - description of 116
 - setting for a migrated file 116
- password 156
 - changing 22
- perform reconciliation tasks 80
- policy information, displaying 29
- premigrated files
 - back up 62
 - deleting from Tivoli Storage Manager storage 79
 - restore 62
- premigrated files database
 - description of 161
 - updating 79
- premigration
 - description of 56
 - explanation of 2
 - starting after threshold migration 57
- premigration percentage 46
 - description of 56
- premigration setting
 - setting 124
- premgrdb.dir file 161
- premgrdb.pag file 161
- prioritizing files for migration 47, 57

Q

- quota
 - description of 47
 - setting 47, 124

R

- reactivating space management 127
 - globally 52, 133
 - on a file system 52
- read-without-recall
 - defined 3

- read-without-recall mode
 - description of 72, 116
 - setting for a migrated file 116
- recall
 - selective 2
 - selective recall 71, 76, 148
 - transparent 71
 - migrate-on-close 2
 - normal 2
 - read-without-recall 2
 - transparent recall 71
- recall daemon
 - description of 83
 - starting 83, 150
- recall ID 147
- recall mode
 - change recall mode 77
 - description of 72, 74
 - dsmattr command 77
 - migrate-on-close 72, 116, 143
 - normal 116, 143
 - partial file recall 74
 - read-without-recall 72, 116
 - set recall mode 77
 - setting for a command 143
 - setting for a migrated file 116
 - streaming recall mode 74
- recall process
 - displaying status of 147
 - removing from queue 153
- recalling files
 - recording information about 161
 - selectively 148
- recalling migrated files
 - automatically 71
 - selectively 76
- recalling your files 71
 - using the dsmattr command 77
 - using the dsrecall command 77
- recalling your premigrated files 2
- reconcileinterval option 28, 82, 110
- reconciliation
 - automatic 79
 - description of 79
 - manual 79
 - starting manually 80, 151
- reconciliation tasks
 - performing 80
- reconciling file systems
 - starting manually 80, 151
- reconciling your file systems
 - automatically 79
 - dsmreconcile command 3
 - recreating deleted stub files 63, 141
- registering with server
 - using closed registration 21
 - using open registration 21
- removing
 - space management from a file system 127
- requirements
 - environment
 - for AIX GPFS 8, 16
 - hardware
 - for AIX GPFS 8
 - for Linux86 GPFS 16

- requirements *(continued)*
 - software
 - for AIX GPFS 8
 - for Linux86 GPFS 16
- requirements, environment
 - for AIX JFS 11
- reset in the SDR 130
- restore migrated files 62
- restore your file systems 67
- restoremigstate option 111
- restoring
 - deleted stub files 63, 141
 - entire system 67
 - file systems 65
 - operating system 67
- restoring files 61
 - migrated files 61
 - premigrated files 61
- rollback 135
- root daemon
 - starting 154
- root user tasks 146
 - adding space management to a file system 44, 124
 - assign management classes 29
 - building a migration candidates list 57, 151
 - copying dsm.opt file 29
 - creating stub files for premigrated files 141
 - deactivating space management 51, 127
 - displaying schedule information 87
 - excluding files from space management 32
 - exporting 128
 - globally deactivating space management 51, 133
 - globally reactivating space management 52, 133
 - importing 128
 - maintaining space management 124
 - modifying 29
 - reactivating space management 52, 127
 - recreating deleted stub files 141
 - removing space management 53, 127
 - selecting space management settings 45
 - setting and changing password 156
 - setting space management options 26
 - setting up central scheduling 87
 - setting up space management 25
 - start threshold migration 57
 - starting a recall daemon 83, 150
 - starting a root daemon 154
 - starting a scout daemon 155
 - starting reconciliation 80, 151
 - starting the client scheduler 87
 - starting the space monitor daemon 83
 - starting threshold migration 118
 - summary of 1
 - synchronizing client and server 81, 151

- root user tasks (*continued*)
 - updating space management settings 50, 124
 - use a default management class 29

S

- sample script
 - start HACMP 37
 - start IBM Cluster 1350 40
 - stop HACMP 37
 - stop IBM Cluster 1350 40
- scheduled services
 - displaying information about completed work 87
 - displaying schedule information 87
- scheduling
 - backup-archive services 87
 - space management services 87
- scout daemon
 - starting 83, 155
- selected files
 - migrating 58, 139
 - recalling 76, 148
- selecting space management settings 45
- selecting your options 25
- selective migration
 - defined 1
 - description of 58
 - performing 58, 139
- selective recall
 - defined 3
 - description of 76
 - performing 76, 148
- setting
 - execution modes 143
 - recall mode for a migrated file 116
- setting environment variables
 - DSM_CONFIG 34
 - DSM_DIR 34
 - DSM_LOG 34
- setting space management options 25
- setup tasks
 - optional 29
- size factor
 - description of 47
 - setting 47, 124
- size of stub files 48
- software requirements
 - for AIX GPFS 8
 - for AIX JFS 11
 - for Linux86 GPFS 16
 - HP-UX 14
 - Solaris 19
- Solaris
 - client components 19
 - communication methods 19
 - considerations for installing HSM 20
 - hardware requirements 19
 - installing 19, 20
 - operating system requirements 19
 - software requirements 19
- space management
 - adding to a file system 44, 124
 - adding to file systems 48
 - adding to nested file systems 49
 - maintaining 124
 - space management (*continued*)
 - removing from a file system 53, 127
 - setting up 25
 - updating a file system 124
 - space management daemons
 - recall 82
 - space monitor 82
 - space management files
 - descriptions of 161
 - space management objects
 - candidates file 161
 - dsmmigfstab file 44
 - logdir directory 161
 - orphan.stubs file 161
 - premigfdb.dir file 161
 - premigfdb.pag file 161
 - status file 161
 - space management options 26
 - checkfororphans 28
 - checkthresholds 28
 - errorprog 28
 - kernelmessages 28
 - maxmigrators 28
 - maxrecalldaemons 28
 - maxreconcileproc 28
 - maxthresholdproc 28
 - migfileexpiration 28
 - migrateserver 28
 - minminrecalldaemons 28
 - reconcileinterval 28
 - space management services
 - scheduling 87
 - space management settings
 - age factor 47, 124
 - descriptions of 45
 - displaying 131
 - high threshold 46, 124
 - low threshold 46, 124
 - maxcandidates 124
 - premigration percentage 46
 - premigration setting 124
 - quota 47, 124
 - selecting 45
 - setting 45, 124
 - size factor 47, 124
 - stub file size 48, 124
 - updating 50, 124
 - space monitor daemon 94
 - description of 83
 - starting 83, 146
 - space usage information, displaying 85, 120, 121
 - spacemgtechnique attribute 29
 - start threshold migration
 - dsmautomig command 57
 - starting
 - client scheduler 87
 - starting automatic reconciliation 79
 - starting reconciliation manually 80
 - starting the HSM client 23
 - status file
 - description of 161
 - updating 79
 - stopping the daemons 84
 - streaming recall mode
 - defined 3
 - description of 74

- stub file
 - defined 1
 - orphan.stub 79
 - orphaned 79
- stub file size
 - choosing 48
 - default size 48
 - description of 48
 - minimum size 48
 - setting 124
- stub files
 - recreating 141
- summary of HSM 1
- synchronizing client and server
 - starting manually 151
- synchronizing client-server
 - description of 79
 - starting manually 80
- synchronizing the client-server
 - starting manually 81

T

- takeover filespec 134
- tasks
 - optional setup 29
- threshold migration 56
 - defined 1
 - setting thresholds for 46
 - start manually 57
 - starting manually 118
- thresholds
 - high threshold 46
 - low threshold 46
- timestamp control mode 143
- Tivoli products and documentation
 - feedback survey xi
 - ordering xi
- Tivoli Storage Manager files
 - for HP-UX 14
 - for Solaris 19
- Tivoli Storage Manager password 156
- transfer HSM management of file system 134, 135
- transparent recall 71
 - defined 2

U

- UNIX clients
 - environment variables, setting 34
- update files in SDR 130
- updating space management settings 50, 124
- using commands 113
- utilities commands
 - dsmdf 85
 - dsmdu 85
 - dsmls 85
 - dsmmigfs 85
 - dsmmigundelete 85

V

- variables
 - Bourne and Korn shell, setting 35

variables (*continued*)
 C shell, setting 35
variables, environment
 DSM_CONFIG 34
 DSM_DIR 34
 DSM_LOG 34
vfs for AIX JFS 13
virtual file system, AIX JFS 13

W

workstation
 registering with server 21



Program Number: 5698-ISM
5698-ISX
5698-SAN

Printed in USA

GC32-0794-02



Spine information:



IBM Tivoli Storage Manager
for Space Management for UNIX

User's Guide

Version 5
Release 2