

IBM Storage Protect HSM for Windows Administration Guide

Version 8.2.0



Contents

| | |
|--|-----------|
| List of Tables | 5 |
| Who should read this publication | 8 |
| Publications | 8 |
| Conventions used in this manual | 8 |
| The following features are new in version 8.1.26: | 8 |
| The following features are new in version 8.1.25: | 8 |
| The following features are new in version 8.1.24: | 8 |
| The following features are new in version 8.1.22: | 8 |
| The following features are new in version 8.1.21: | 9 |
| The following features are new in version 8.1.20: | 9 |
| The following features are new in version 8.1.19: | 9 |
| The following features are new in version 8.1.17: | 9 |
| The following features are new in version 8.1.15: | 9 |
| The following features are new in version 8.1.13: | 9 |
| The following features are new in version 8.1.12: | 9 |
| The following features are new in version 8.1.11: | 9 |
| The following features are new in version 8.1.10: | 9 |
| The following features are new in version 8.1.9: | 10 |
| The following features are new in version 8.1.7: | 10 |
| The following features are new in version 8.1.4: | 10 |
| The following features are new in version 8.1.2: | 10 |
| The following feature is new in version 8.1: | 11 |
| HSM for Windows™ client overview | 12 |
| Migration overview | 13 |
| Migration types | 13 |
| Recall modes | 15 |
| Stub files | 15 |
| Previously migrated files | 16 |
| Retention of migrated files in IBM Storage Protect™ storage | 17 |
| Reconciliation overview | 17 |
| Client commands and GUI overview | 17 |
| Getting started with HSM for Windows™ client web GUI | 18 |
| Login page | 18 |
| Dashboard page | 18 |
| Activity page | 18 |
| Alerts page | 19 |
| Configuration page | 19 |
| Support page | 21 |
| HSM for Windows dashboard overview | 21 |
| Tiles description | 21 |
| Purpose and usage | 22 |
| Installation of the HSM for Windows™ client | 23 |
| Planning to install the HSM for Windows™ client | 23 |
| Hardware and software requirements | 23 |
| National language environments | 23 |
| Compatibility with other software | 23 |
| Restrictions for rolling back | 24 |
| Preparing for the installation | 25 |
| Installing the HSM for Windows™ client | 25 |
| Installing and configuring the HSM for Windows™ client in a cluster environment | 25 |
| HSM in cluster environments | 26 |
| Upgrading the HSM for Windows™ client | 28 |
| Migrating Windows™ alternate data stream data for files that were migrated before version 7.1.2 | 28 |
| Configuring the HSM for Windows® client | 29 |
| Configuring the connection between the HSM for Windows™ client and the IBM® Storage Protect server | 29 |
| Client password-character restrictions | 31 |
| Configuring the HSM client to connect to a secondary IBM® Storage Protect server | 32 |
| Configuring the retention period of migration copies | 32 |
| Changing the retention period of migration copies | 34 |
| Configuring a new file space | 34 |

| | |
|--|-----------|
| Configuring regional settings..... | 35 |
| Excluding Windows™ alternate data stream names..... | 35 |
| HSM advanced parameters and preferences settings | 35 |
| File location preferences..... | 36 |
| Move Settings | 37 |
| File recall quotas..... | 37 |
| Recall service settings..... | 41 |
| Tracing preferences..... | 41 |
| Managing space with HSM for Windows® | 43 |
| Migration jobs..... | 43 |
| Creating migration jobs | 43 |
| Examples of including and excluding files | 45 |
| File groups..... | 47 |
| Calculating a migration job's space savings | 48 |
| Migration jobs start by a schedule, GUI, or CLI | 48 |
| Removing unused stubs from a file system..... | 49 |
| Migration by file list..... | 51 |
| Threshold migration..... | 51 |
| Migration candidates | 51 |
| Migration triggers..... | 52 |
| Configuring threshold migration..... | 52 |
| Space management of the system volume | 55 |
| Selectively retrieving and recalling migrated files..... | 55 |
| Retrieving migrated files | 56 |
| Selectively recalling migrated files..... | 57 |
| Automatic backup before migration | 57 |
| Choosing a backup options file | 58 |
| Backup and restore of migrated files..... | 58 |
| Options for backing up of migrated files | 59 |
| Managing backup performance when stub file encryption changes | 60 |
| Backing up migrated files separately from resident files | 61 |
| Options for restoring migrated files..... | 61 |
| Reconciliation..... | 63 |
| Changed volume mount-paths | 64 |
| Configuring reconciliation with the graphical user interface..... | 65 |
| Space requirements for reconciliation | 66 |
| Previewing files that would be deleted by a reconciliation process | 66 |
| Deleting protected files from IBM® Storage Protect storage | 66 |
| Moving migrated files..... | 67 |
| Migrated data is automatically moved when stub files are moved..... | 67 |
| Stub files in moving state | 68 |
| Moving stub files to another location | 68 |
| Continuing HSM services when a volume or file server is renamed | 69 |
| Renamed volume drive letters or mount points..... | 69 |
| New file server hardware or a changed file server host name..... | 70 |
| Hardware volume mappings..... | 70 |
| Mapping volumes..... | 70 |
| Displaying HSM listing files..... | 71 |
| HSM for Windows™ client REST API and web client | 72 |
| Monitoring HSM-managed volumes with Prometheus and Grafana open platforms..... | 72 |
| Configuring the HSM for Windows™ client web GUI..... | 73 |
| HSM for Windows® commands | 75 |
| Entering command parameters | 75 |
| Restriction for running a command again | 75 |
| Command parameters override default and job settings | 75 |
| Client return codes..... | 76 |
| dsmclc.exe | 76 |
| dsmclc createfilespace | 76 |
| dsmclc defaults | 77 |
| dsmclc delete | 78 |
| dsmclc legend | 80 |
| dsmclc list | 80 |
| dsmclc listfilespace s..... | 82 |
| dsmclc listmgmtclasses | 84 |
| dsmclc migrate | 85 |
| dsmclc migratelist | 87 |
| dsmclc recall | 89 |
| dsmclc recalllist | 90 |
| dsmclc register | 92 |
| dsmclc retrieve | 94 |
| dsmfileinfo.exe | 96 |
| Syntax | 96 |
| Options | 96 |

| | |
|---|------------|
| Examples | 97 |
| dsmfind.exe | 97 |
| Syntax | 98 |
| Options | 98 |
| Examples | 99 |
| dsmhsmclc.exe | 99 |
| Managing reconciliation tasks | 99 |
| Managing threshold migration | 104 |
| Managing scans for volume statistics | 108 |
| dsminfo.exe | 110 |
| Syntax | 110 |
| Options | 110 |
| Examples | 111 |
| dsmmove.exe | 111 |
| Syntax | 111 |
| Parameters | 111 |
| Examples | 113 |
| dsmquota.exe | 113 |
| Syntax | 114 |
| Parameters | 114 |
| Example | 115 |
| dsmtool.exe | 115 |
| Syntax | 116 |
| Parameters | 116 |
| Example | 116 |
| Troubleshooting the HSM for Windows™ client | 117 |
| Troubleshooting steps and information | 117 |
| Trying the action again | 117 |
| Collecting data and files for IBM® support | 117 |
| Offline stub files are recalled when they are first synchronized | 117 |
| Problems with VSS during reconciliation | 117 |
| Small migrated files occupy much space on IBM® Storage Protect server storage | 118 |
| Small file size for threshold migration might require a cleanup | 118 |
| Accessibility features for the IBM® Storage Protect product family | 119 |
| Overview | 119 |
| Keyboard navigation | 119 |
| Interface information | 119 |
| Vendor software | 119 |
| Related accessibility information | 119 |
| Notices | 120 |
| Trademarks | 121 |
| Terms and conditions for product documentation | 121 |
| Privacy policy considerations | 122 |
| Glossary | 123 |
| Index | 124 |

List of Tables

| | |
|--|-----|
| Table 1 | 8 |
| Table 2: Migration jobs compared to threshold migration | 14 |
| Table 3: Advanced parameters descriptions | 36 |
| Table 4: Tracing preferences: Trace levels | 41 |
| Table 5: Tracing preferences: Trace file size | 42 |
| Table 6: Tracing preferences: Log file size settings | 42 |
| Table 7: Example base file set | 45 |
| Table 8: Interaction of options skipmigrated and checkreparsecontent during incremental backup | 60 |
| Table 9: Results of using restoremigstateand restorecheckstubaccessoptions. | 62 |
| Table 10: HSM for Windows™ client Command Prompt window commands..... | 75 |
| Table 11: An explanation of client return codes | 76 |
| Table 12: Options for dsmfileinfo.exe | 97 |
| Table 13: Options for dsminfo.exe..... | 110 |

Note:

Before you use this information and the product it supports, read the information in “Notices” on page 120.

This edition applies to version 8, release 2 of IBM® Storage Protect HSM for Windows™ (product number 5725-X14), and to all subsequent releases and modifications until otherwise indicated in new editions.

About this publication

This publication provides the information to install, configure, monitor, and troubleshoot problems with the IBM® Storage Protect HSM for Windows™.

Who should read this publication

This publication is intended for persons who are responsible for installing, configuring, monitoring, and troubleshooting the IBM® Storage Protect HSM for Windows™. In this publication, it is assumed that you have a working knowledge of the IBM® Storage Protect HSM for Windows™.

Publications

The IBM® Storage Protect product family includes IBM® Storage Protect Plus, IBM® Storage Protect for Virtual Environments, IBM® Storage Protect for Databases, and several other storage management products from IBM®.

To view IBM® product documentation, see [IBM® Documentation](#).

Conventions used in this manual

This manual uses the following typographical conventions:

| Example | Description |
|--------------------|---|
| cancel | Boldface type indicates a parameter or a user interface control. |
| <i>optionvalue</i> | Italic type indicates a placeholder for information you provide, or for special emphasis in the text. |
| user input | Monospace type indicates fragments of a program or information as it might appear on a display screen, such as a command example. |
| plus sign (+) | A plus sign between two keys indicates that you press both keys at the same time. |

New for IBM® Storage Protect HSM for Windows™

Learn about new features and updates for IBM® Storage Protect HSM for Windows™ version 8.2.0

The following features are new in version 8.1.26:

The web-client displays the latest alerts on the dashboard page until it is reviewed.
Any new or updated content is indicated by a vertical bar (|) in the margin.

The following features are new in version 8.1.25:

Maintenance updates are provided in this release.

The following features are new in version 8.1.24:

HSM for Windows™ client web client

Beginning with version 8.1.24, the HSM for Windows™ client web GUI now supports configuration of log file settings for different services.

The HSM for Windows™ client web GUI now supports OpenSSL 3.0, 3.1, 3.2, and 3.3.

The following features are new in version 8.1.22:

HSM for Windows™ client web client

Beginning with version 8.1.22, the IBM® Storage Protect HSM for Windows™ client web GUI now supports a per client IBM® Storage Protect server status on the installation page of the web client. The dashboard page of the web client displays the summarized status in a chart diagram with details in the diagram's tooltip.

The following features are new in version 8.1.21:

HSM for Windows™ client web client

Beginning with version 8.1.21, the IBM® Storage Protect HSM for Windows™ client web GUI now supports different security levels for encryption strength in the communication.

The following features are new in version 8.1.20:

Beginning with version 8.1.20, the product has been fully rebranded as IBM® Storage Protect HSM for Windows™. The HSM for Windows™ client web GUI now supports OpenSSL 3.0/3.1.

In this release, an APAR was fixed.

The following features are new in version 8.1.19:

Beginning with version 8.1.19, IBM Storage® Protect HSM for Windows™ is now IBM® Storage Protect HSM for Windows™. The IBM Spectrum® family of products has been rebranded as IBM® Storage, optimizing performance for data-intensive applications with intelligent edge, core, and cloud storage solutions. To learn more about the rebranding transition, see [IBM Spectrum® brand change to IBM® Storage](#).

The following features are new in version 8.1.17:

HSM for Windows™ client web client

Beginning with version 8.1.17, the IBM® Storage Protect HSM for Windows™ client web GUI allows administrators to create, delete, and configure migration jobs.

The following features are new in version 8.1.15:

HSM for Windows™ client web client

Beginning with version 8.1.15, the IBM® Storage Protect HSM for Windows™ client web GUI allows administrators to start and stop migration jobs and displays migration configurations.

The following features are new in version 8.1.13:

HSM for Windows™ client web client

Beginning with version 8.1.13, the IBM® Storage Protect HSM for Windows™ client web GUI allows administrators to start and view volume scans, create filespace, delete alerts, and reset user quotas.

In this release, an APAR was fixed.

The following features are new in version 8.1.12:

HSM for Windows™ client web client

Beginning with version 8.1.12, the IBM® Storage Protect HSM for Windows™ client web GUI allows administrators to change application configuration and download support logs. The volume scan feature has been also enhanced.

The following features are new in version 8.1.11:

HSM for Windows™ client web client

In version 8.1.11, a new HSM for Windows™ client web client is now available. With the HSM for Windows™ client web client, administrators can use a web browser to log in to the HSM for Windows™ client web client, access one of the HSM nodes in the company's installation, and supervise all the HSM-managed file servers. For more information, see [“HSM for Windows client REST API and web client” on page 72](#) and [“Configuring the HSM for Windows client web GUI” on page 73](#).

The following features are new in version 8.1.10:

REST API

In version 8.1.10, the HSM for Windows™ client REST API functionality for the Prometheus time-series database is extended with many metrics that are collected during HSM for Windows™ client activity. The Prometheus open source system can connect to HSM-managed file servers, monitor free volume space,

migration and recall activity, and backup and dsmmove activity. In addition, an example HSM dashboard for the Grafana Analytics Platform is now also provided.
For a complete list of the metrics and more information, see [“Monitoring HSM-managed volumes with Prometheus and Grafana open platforms” on page 72.](#)

The following features are new in version 8.1.9:

REST API

In version 8.1.9, the HSM for Windows™ client provides REST API functionality to access metrics through a web application. The **hsmnet.exe** service provides recall and migration counts and volume space data. For example, the Prometheus open-source system can connect to HSM-managed file servers, and to monitor recall and migration activity, and free space.

For more information, see [“Monitoring HSM-managed volumes with Prometheus and Grafana open platforms” on page 72.](#)

Volume statistics

Volume statistics, like the number of stubs, the logical size of all stubs, and other statistics on a HSM-managed volume can be collected in a scan for volume statistics. The scan for volume statistics can be configured from the command prompt using the **dsmhsmc1c.exe** command.

For more information, see [“Managing scans for volume statistics” on page 108.](#)

The following features are new in version 8.1.7:

Recall enhancements

In version 8.1.7, the recall operation is now much faster and will not time out. By using Server Message Block (SMB), which allows systems within the same network to share files, migrated files start copying immediately, and you do not need to wait until the recall operation is complete. If the application that opens the file uses this feature, it can read or display the content of the file while the remaining parts of the file are being recalled.

For more information, see [“Recall modes” on page 15.](#)

The following features are new in version 8.1.4:

Threshold Migration processing now includes aBinFileLimitPercentoption

The **dsmhsmc1c.exe** command, used to configure threshold migration from the **Command Prompt** window, now includes aBINFILELIMITPERCENToption. This option is used to configure the maximum size of a candidates list file as a percentage of the disk size.

For more information, see [“Managing threshold migration” on page 104.](#)

Filter driver performance enhancements

In version 8.1.4, a transparent recall of a migrated file is now much faster. Under normal conditions, the recall duration is 50 % or less compared to versions before version 8.1.4.

For more information, see [“Recall modes” on page 15.](#)

The following features are new in version 8.1.2:

Result filters now supported in the Command Line Interface

In version 8.1.0, the graphical user interface (GUI) was enhanced to include a **Result filters** tab, enabling you to further refine your search results when you retrieve migrated files. The filters included **Migration action** and **Migration time**. In version 8.1.2 these result filters are now also supported on the Command Line Interface.

For instructions about retrieving migrated files, see [“Retrieving migrated files” on page 56.](#)

Filter driver performance enhancements

In addition, version 8.1.2 also includes recall performance enhancements to handle more parallel requests, while also using less system resources.

Password Promptoption no longer supported

The IBM® Storage Protect HSM for Windows™ **Client Configuration Wizard** no longer supports the Password Promptoption.

For more information, see [“Configuring the connection between the HSM for Windows client and the IBM Storage Protect server” on page 29.](#)

The following feature is new in version 8.1:

Refine the search results when you retrieve migrated files

The version 8.1 graphical user interface (GUI) is enhanced to include a **Result filters** tab, enabling you to further refine your search results when you retrieve migrated files. The filters include **Migration action** and **Migration time**:

- Migration action is the action that is specified to the file found on the backend server. The migration action of files that are migrated with HSM versions earlier than version 7.1.4 is unknown.
- Migration time is the time when the backend file was last migrated. The migration time of files that are migrated with HSM versions earlier than version 7.1.6 is unknown.

For instructions about retrieving migrated files, see [“Retrieving migrated files” on page 56](#).

HSM for Windows™ client overview

The IBM® Storage Protect HSM for Windows™ client provides hierarchical storage management (HSM) for Windows® New Technology File System (NTFS) and Resilient File System (ReFS) file systems.

The figure shows an overview of hierarchical storage management.

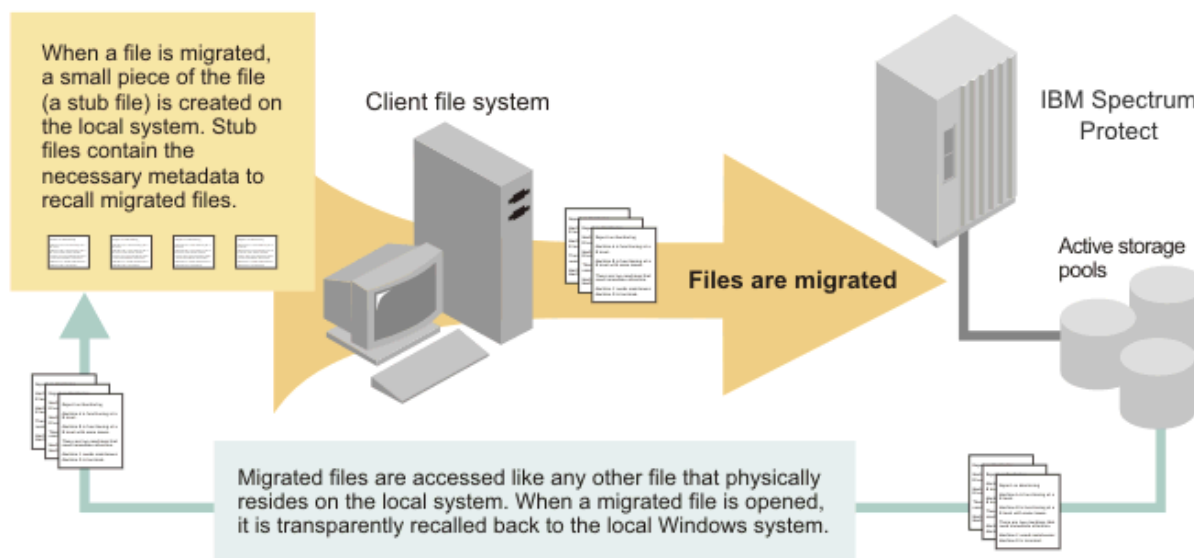


Figure 1: Overview of hierarchical storage management

HSM is a data storage system that automatically moves data between high-cost and low-cost storage media. HSM exists because high-speed storage devices, such as hard disk drives, are more expensive per byte stored than slower devices, such as magnetic tape drives. You can use HSM to store the bulk of your enterprise's data on slower devices, and then copy data to faster disk drives only when needed.

In effect, HSM turns the fast disk drives into caches for the slower mass storage devices. The HSM for Windows™ client monitors the way files are used and automates policies for migrating files to slower devices.

The HSM for Windows™ client manages the migration of individual files, files from parts of file systems, or complete file systems, to remote storage in IBM® Storage Protect. Migrated files can be accessed, opened, and updated by the Windows® application corresponding to the file extension.

In addition to the migration and recall of files and the reconciliation of file systems, the HSM for Windows™ client provides extra functions beyond the scope of traditional HSM:

- An administrator can define migration jobs for each volume. The job can include or exclude files of a certain file type (extension). Files can be included or excluded depending on file age or size. The files that are eligible for each migration job can be stored in separate file spaces in IBM Storage Protect™ storage.
- An administrator can define recall quotas to limit the number of file recalls during a specified time period. Quotas can apply to the entire system, to user groups, or to specific users.
- The HSM for Windows™ client can also be used for archiving purposes. In this case, files are migrated to IBM Storage Protect™ storage and the original files are either kept on disk or deleted.
- Search and retrieve options are available to the administrator for migrated files. Selected files or complete file spaces can be retrieved either to their original location in the file system or to a different location in the file system.
- When migrated files are recalled and changed by a user, several versions of a migrated file are kept in IBM Storage Protect™ storage until the file system is reconciled. A user recall always accesses the latest version of a file. However, an administrator can retrieve any available version of a file.
- Threshold migration monitors file-system space usage and migrates files when space is needed.
- Threshold migration migrates older and larger files from your file system. You configure whether file age or file size is a better qualifier for migration.
- You can move migrated data without interrupting HSM services.

You can move migrated files to accommodate the changing needs of users, applications, and hardware. For example, if a user moves to another site, you can move the migrated data. If a new or changed application requires that data is moved to another location, you can move the migrated files. You can maintain HSM services without recalling and migrating the files again.

- You can replace or rename a volume or file server and maintain HSM services without recalling and migrating the files again.

The following are some advantages beyond the classical HSM approach:

- The scope of individual migration jobs can be limited by the number of files and data volume.
- Individual jobs can be run at different times.
- Migration jobs can be organized according to the logical structure of a volume (including different parts of the directory structure). Jobs can reflect the structure of an organization or user groups.
- Migration jobs can be organized according to different types of files such as office documents, images, and text files. This organization provides a more logical view on data than traditional HSM.
- Threshold migration can automatically prevent your volumes from running out of free space.
- With threshold migration's age weighting, active files are kept on the volume. Less-active files are migrated to IBM Storage Protect™ storage.
- With threshold migration's size weighting, larger files are migrated to IBM Storage Protect™ storage. Larger files provide a more efficient migration.
- You can implement migration jobs and threshold migration on the same volume. You can build a policy that is based on both file values (migration jobs) and space usage (threshold migration).

The HSM for Windows™ client comes with a graphical user interface (HSM for Windows™ client GUI). You can use the HSM for Windows™ client GUI to define and run migration jobs, threshold migration, reconciliation, searches and file retrieval, and to define general settings. You can also do many of these tasks by using HSM for Windows™ client commands from a **Command Prompt** window.

The HSM for Windows™ client supports local, fixed NTFS and ReFS file systems. Microsoft® Cluster Server (MSCS) cluster volumes are supported, if they are formatted in NTFS or ReFS. Windows® File Allocation Table (FAT) partitions, Common Internet File System (CIFS) shared folders, network-attached storage (NAS) drives, and other file systems are not supported.

Migration overview

A migration process copies files from a file system to IBM® Storage Protect storage. The migrated copies are returned to the file system as required.

There are several ways to migrate files to IBM Storage Protect™ storage, and several ways to get the files back to the file system.

Migration types

You can configure migration jobs and threshold migration. You can selectively migrate files that are specified in a list file.

Migration jobs and list migrations specify which files can be migrated, but they do not consider the space capacity of the volume. Threshold migration controls space usage of the volume, but allows less control of which files are migrated.

Migration jobs

A migration job defines a set of files and their migration behavior. When you run the job, the files that are specified in the job are copied to IBM® Storage Protect storage.

A migration job can replace the original file with a stub file, delete the original file, or do nothing to the original file. You configure the action. You configure whether files are backed up before migration.

You can start the migration job immediately with the HSM for Windows™ client GUI or with an HSM for Windows™ client command from a **Command Prompt** window. You can also start the migration job later with a scheduling program acquired from another vendor.

List migrations

A list migration migrates the files that are listed in a text file. A list migration is not affected by disk-space usage or the age and size of files.

A list migration can replace the original file with a stub file, delete the original file, or do nothing to the original file. You configure the action. You configure whether files are backed up before migration.

Start a list migration with the HSM for Windows™ client **dsmclic migratelist** command.

Threshold migration

Threshold migration provides migration that is based on space usage. When the used space on a volume reaches a high threshold, migration begins automatically. Files are migrated to free up space until used space falls to a low threshold. The files that are migrated meet a minimum age and size, and are prioritized for migration. Less dynamic and larger files are migrated before more dynamic and smaller files. With proper configuration, threshold migration can automatically prevent the volume from running out of space. Threshold migration replaces the original file with a stub file. You configure whether files are backed up before migration.

Configure threshold migration with the HSM for Windows™ client **dsmhsmclic configurethresholdmig** command.

The following table summarizes the similarities and differences between migration jobs, list migrations, and threshold migration.

The table is a summary of differences and similarities between migration jobs, list migrations, and threshold migration.

| Table 2: Migration jobs compared to threshold migration | | |
|--|--|--|
| Criterion | Migration job and list migration | Threshold migration |
| Which files are migrated? | <p>Migration job:</p> <p>You configure the path, type (file extension), minimum age, and minimum size of files to migrate. All files that meet the criteria are migrated.</p> <p>List migration:</p> <p>The files are identified in a list file.</p> | You configure the minimum file age and minimum file size, and the importance of file age relative to file size. HSM for Windows™ client creates a ranked list of migration candidates that are based on the criteria. Files from this list are migrated as needed to meet the space usage targets. |
| When does migration occur? | You start migration manually, or with a scheduling tool that is provided by another vendor. | HSM for Windows™ client automatically starts migration when it detects that space usage on the volume reaches the high threshold. |
| When does migration end? | <p>Migration job:</p> <p>Migration ends when all files that meet the criteria are migrated.</p> <p>List migration:</p> <p>Migration ends when all files in the list are migrated.</p> | Migration ends when space usage on the volume reaches the low threshold, or when there are no more candidates for migration. |
| What remains on the volume from which the files were migrated? | <p>HSM for Windows™ client can do one of three things, as you configure:</p> <ul style="list-style-type: none"> Replace the original file with a stub file Leave the original file Delete the original file, create no stub file. | HSM for Windows™ client replaces the original file with a stub file. |

| Criterion | Migration job and list migration | Threshold migration |
|---|---|---|
| When are files automatically recalled to the originating file system? | If the file system requests an operation that cannot be satisfied by the stub file, the migrated file is automatically and transparently recalled. The stub file provides the information to recall the file. | If the file system requests an operation that cannot be satisfied by the stub file, the migrated file is automatically and transparently recalled. The stub file provides the information to recall the file. |
| Can I selectively retrieve the migrated files? | Yes, by using the HSM for Windows™ client GUI or the dsmc1c retrieve command. | Yes, by using the HSM for Windows™ client GUI or the dsmc1c retrieve command. |
| Can I selectively recall the migrated files? | Yes, if a stub file exists on the file system. Use the dsmc1c recall command or the dsmc1c recalllist command. | Yes, if a stub file exists on the file system. Use the dsmc1c recall command or the dsmc1c recalllist command. |

Recall modes

Migrated files can be recalled transparently, recalled selectively, and retrieved selectively.

A file is recalled automatically when you or a Windows™ application accesses the stub file. You can manually return a migrated file to the file system by using the information on the IBM Storage Protect™ server or the information in stub files.

Transparent recall

When you or a Windows™ application accesses a migrated file stub, the HSM for Windows™ client automatically recalls the file from IBM Storage Protect™ storage.

If only the Windows™ alternate data stream (ADS) data in a stub file is accessed or modified, the file is not recalled. When a file is recalled because the primary data stream (PDS) data is accessed in the stub file, the ADS data is not recalled. The ADS data is stored in the stub file on the file system and does not change when the PDS data is recalled.

Transparent recall operations are now much faster and will not time out. By using Server Message Block (SMB), which allows systems within the same network to share files, migrated files start copying immediately, and you do not need to wait until the recall operation is complete. If the application that opens the file uses this feature, it can read or display the content of the file while the remaining parts of the file are being recalled. This is especially beneficial when recalling very large migrated files, as by using SMB the recall operation will not time out.

Selective recall

You can selectively recall migrated files that were replaced with stubs when they were migrated. You can search the file system for stub files that match a pattern. You can recall the migrated files that are listed in a text file. The files in the text file must be stub files. Security attributes and ADS data are not recalled with the file.

Selective retrieve

You can selectively retrieve migrated files by using information from the IBM Storage Protect™ server. You can specify whether security attributes and ADS data are retrieved with the file. When you retrieve migrated files, stub files are not required.

Stub files

A stub file is created on the file system from which a file is migrated. The stub file contains information for the HSM for Windows™ client to recall the original file to the file system.

A stub file contains the same recall information whether it was created by a migration job, a list migration, or a threshold migration.

When you or a Windows™ application accesses a migrated file stub, the HSM for Windows™ client automatically recalls the file from IBM Storage Protect™ storage. This automatic recall is called transparent recall.

Restriction: If only the Windows™ alternate data stream (ADS) data in a stub file is accessed or modified, the file is not recalled. When a file is recalled because the primary data stream (PDS) data is accessed in the stub file, the ADS data is not recalled. The ADS data is stored in the stub file on the file system and does not change when the PDS data is recalled.

A stub file looks and acts like a regular file on the file system, with a few exceptions:

- Files that are migrated are marked.
 - In the **Windows Explorer**, a migrated file has an overlay icon.
 - On a **Command Prompt** window, a migrated file is enclosed in brackets.
- Access to migrated files can be slower, if the file operation recalls the migrated file from IBM® Storage Protect storage.

When a file is migrated, the last access time of the file does not change.

You can selectively recall a file from IBM Storage Protect™ storage if a stub file exists on the file system.

Previously migrated files

After a file is migrated, it can be migrated again. Whether the file is migrated again depends on how the file was last migrated and how the file is changed.

When a file is recalled, modified, and migrated again, that new version of the file is stored in IBM® Storage Protect storage. More than one version of the file exists in IBM Storage Protect™ storage until the file system is reconciled. Any file operation that requires the file to be recalled yields the most recently migrated version.

If a file was migrated and replaced with a stub, the subsequent migration of the file depends on how the file changes and the type of migration. A threshold migration does not migrate a stub. For migration types other than threshold migration, the subsequent migration of the stub file depends on how the file changes.

Data: Only the Windows™ alternate data stream (ADS) data changes

The file content changes

When you change the content of a file, the HSM for Windows™ client recalls the file from IBM Storage Protect™ storage. The next time that the file is the object of a migration job or a threshold migration, the new version of the file is migrated. The IBM Storage Protect™ server maintains versions of the migrated file until you run reconciliation. The migrated file is bound to the management class that is specified by the last migration job or threshold migration.

Only the file attributes or times (creation time or last modification time) change

When only file attributes or times (creation time or last modification time) change, the file is not migrated again to IBM Storage Protect™ storage. Instead, the attributes or file times are updated in the IBM Storage Protect™ metadata database the next time. The updates are made the next time that the file is the object of a migration job. The management class does not change, even if the migration job specifies a different management class.

Only the file security attributes change

The second migration of the file depends on whether you configured migration of file security attributes.

If you configured migration of file security attributes:

If only the ACL changes, the file is migrated the next time that the file is the object of a migration job or list migration. At the next job or list migration, the file is temporarily recalled, then migrated with the updated ACL. The version number of the file that is tracked by IBM Storage Protect™ does not change. The previous copy of the file in IBM Storage Protect™ storage is deleted. The migrated file is bound to the management class that is specified by the last migration job or threshold migration. If the file is targeted by a threshold migration, the file is not migrated again.

If you configured no migration of file security attributes:

If only the ACL changes, the file is not migrated the next time that the file is the object of a migration job or list migration.

Only the Windows™ alternate data stream (ADS) data changes

The second migration of the file depends on whether you configured migration of ADS data.

If you configured migration of ADS data:

If only the ADS data changes, the file is migrated the next time that the file is the object of a migration job or list migration. At the next job or list migration, the file is temporarily recalled, then migrated with the updated ADS data. The version number of the file that is tracked by IBM Storage Protect™ does not

change. The previous copy of the file in IBM Storage Protect™ storage is deleted. The migrated file is bound to the management class that is specified by the last migration job or threshold migration. If the file is targeted by a threshold migration, the file is not migrated again.

If you configured no migration of ADS data:

If only the ADS data changes, the file is not migrated the next time that the file is the object of a migration job or list migration.

No changes

Even if a migrated file did not change in any way, it can be migrated again. For example, you can migrate a file; recall the file; and not change the file. When you migrate the file again, the HSM for Windows™ client replaces the existing file with a stub that points to the existing file copy in IBM Storage Protect™ storage. The management class does not change, even if the migration job specifies a different management class.

You can configure migration with the `keepoption`. The file is migrated to IBM Storage Protect™ storage, but the file is not replaced by a stub file. When you change the file, the HSM for Windows™ client does not automatically recall the file or track the changes to the file on the IBM Storage Protect™ server. The file remains unchanged on IBM Storage Protect™ storage. If you migrate the file again, the file is bound to the management class that is specified by the last migration job or list migration.

Retention of migrated files in IBM Storage Protect™ storage

A migrated file is stored in IBM® Storage Protect storage and managed as an archive copy group.

Attention: The default setting for management classes deletes migrated files from IBM® Storage Protect storage after 365 days. The files are deleted from storage whether the original file is replaced with a stub, is deleted, or remains on the file system. To store files longer than 365 days, specify a management class that is suitable for retaining the migration copies. Or change the retention period of the default management class. See [“Configuring the retention period of migration copies”](#) on page 32.

Reconciliation overview

Reconciliation is the process of synchronizing a file system with the IBM® Storage Protect server. After the reconciliation cycle completes, exactly one migrated object exists on the IBM Storage Protect™ server for each migrated file.

By removing obsolete objects from IBM Storage Protect™ server storage, reconciliation helps you to reduce your storage and license expenses. Reconciliation also checks that there is a migrated object on the IBM Storage Protect™ server for every stub file on the volume.

The HSM for Windows™ client reconciles automatically at intervals that are specified with the **reconcileinterval** option you define with the HSM for Windows™ client GUI or with the **dsmhsmc1c.exe** command. An administrative user can also start reconciliation manually at any time.

Client commands and GUI overview

After you install and register the HSM for Windows™ client, you can use the HSM for Windows™ client GUI (graphical user interface) or run commands from a Command Prompt window.

Start the GUI with the **dsmgui.exe** executable file in the installation directory. After the GUI is started, you can configure, monitor, and administer space management with the controls in the GUI. You can complete all HSM operations with the GUI, but not all operations are supported by the commands.

You must start the HSM for Windows™ client GUI with administrative rights on the file server on which it is administered. Each file server on which the HSM for Windows® client is installed must be administered locally.

Many operations that you complete with the HSM for Windows™ client GUI, you can also complete with commands from a **Command Prompt** window. Each command has its own executable file, also in the installation directory.

Getting started with HSM for Windows™ client web GUI

HSM for Windows™ client web GUI is a single management UI that allows administrators to manage configuration of all the Windows HSM file servers and associated storage space and to receive warning and error message alerts.

Administrators can use a web browser to log in to the HSM for Windows™ client web GUI, access one of the HSM nodes in the company's installation, and supervise all the HSM-managed file servers.

The HSM for Windows™ client web GUI is an optional, installable component that comes with the HSM for Windows™ client installer package. The HSM for Windows™ client web GUI is installed on one HSM node that acts as the proxy-server node. All other HSM nodes can act as proxy-client nodes and require registration at the proxy-server node.

The HSM for Windows™ client web GUI consists of the following pages:

Login page

The **Login** page allows administrators to login to HSM for Windows™ client web GUI by using valid credentials and access one of the HSM nodes in the company's installation and supervise all the HSM-managed file servers.

Dashboard page

The **Dashboard** page displays information about the HSM nodes in the company's installation. The **Dashboard** page consists the following tiles:

- **HSM Clients:** The tile displays the number of HSM clients connected with the application with status and risk information. The HSM client risk is calculated based on volume space usage and recall capabilities.
- **File Server:** The tile displays the connected file server summary information with used space, free space, migrated space, and migrated files. Migrated space and migrated files values are calculated based on HSM volume scan data.
- **Volumes:** The tile displays the total number of connected volumes along with their risk information.
- **Services:** The tile displays the total number of running services along with their category and status.
- **Alerts:** The tile displays the total number of alert messages at the top. The alert messages are categorized in the Recall, Recall Quota, Migration, Reconcile, and Other classes. Unread alert message counts for each class are highlighted in the blue color.
- **Maintenance:** The tile displays running maintenance tasks information. The maintenance tasks are classified in the Threshold, Volume Scan, and Reconcile types.
- **SP Servers:** The tile displays information about connected IBM® Storage Protect servers.
- **Management Classes:** The tile displays information about management classes.
- **Filespaces:** The tile displays information about file spaces in use.

Activity page

The **Activity** page allows administrators to run and cancel migration jobs, view job migration status, details of inspected and migrated files, and freed space. The **Activity** page consists of the following tabs:

- **Volumes:** Click the **Volumes** tab to view managed volumes on the different client nodes and the status of threshold migration settings. On the **Volumes** page, you can view the information in a tabular format with the following columns.
 - **Fileserver:** The column displays the name of the file server.
 - **Volume:** The column displays the name of the volume on a file server.
 - **Capacity:** The column displays the used and available capacity for each volume. The column is highlighted in green, yellow, or red color. If less than 80 percent capacity is used, column is highlighted in the green color, if used capacity is in between 80 to 90 percent, the column is highlighted in the yellow color, if more than 90 percent capacity is used, the column is highlighted in the red color.
 - **Threshold:** The column displays whether the HSM threshold is configured for a volume.
- **Job Migration:** Click the **Job Migration** tab to view configured migration jobs and run migration tasks. On the **Job Migration** page, you can view the information in a tabular format with the following columns:

- **File Server:** The column specifies the file server name for which a migration job is configured.
- **Job:** The column specifies the name of a migration job.
- **Last Start:** The column specifies the start date of a migration job.
- **Duration:** The column specifies the runtime of a migration job.
- **Inspected Files:** The column specifies the number of inspected files by a migration job.
- **Migrated Files:** The column specifies the number of migrated files by a migration job.
- **Freed Space:** The column specifies space that is freed by a migration job.
- **Result:** The column specifies whether a migration job is completed successfully.
- **Migration Task:** The option allows administrators to run a migration job.
- **Volume Analysis:** Click the **Volume Analysis** tab to view volume scans for each client node. On the **Volume Analysis** page, you can view the information in a tabular format with the following columns:
 - **File Server:** The column specifies the name of a file server.
 - **Volume:** The column specifies a volume on a file server.
 - **Capacity:** The column specifies the total storage capacity of a volume.
 - **Offline:** The column specifies the offline storage space of a volume.
 - **Resident Files:** The column specifies the number of resident files on a volume at the time of last scan.
 - **Stub Files:** The column specifies the number of stub files on a volume at the time of last scan.
 - **% Stub Files:** The column specifies the percentage of stub files on a volume at the time of last scan.
 - **Scan Task:** The column allows administrators to start scan task manually and displays the duration from the last scan.
 - **Expand Icon:** The column allows administrators to view the trends of volume occupancy based on the occupied volume along with time. Click the expand icon to view volume scan graphs of resident and migrated data and files.

Alerts page

The **Alert** page displays all error and warning log messages in tabular format with the most recent message appears at the top. For each alert or warning message, administrators can view the time of alert, server, process, alert class, and an alert message. Click a row to view the detailed error message, explanation of the alert message, system action, and suggested user action. At the top of the table, the number of alert messages for each class is listed. The alert message classes are User Recall, Recall Quota, Migration, Reconcile, and Other.

Configuration page

The **Configuration** page allows administrators to configure and monitor migration jobs. The **Configuration** page consists of the following tabs:

- **HSM Installations:** Click the **HSM Installations** tab to view configured client nodes information. You can view connection status and start and stop HSM services on different nodes. The **HSM Installations** page displays information in the tabular format with the following columns:
 - **File Server:** The column specifies the name of a file server.
 - **Cluster Name:** The column specifies the name of a cluster.
 - **HSM Version:** The column specifies the HSM version that is installed on a file server.
 - **SP Node:** The column specifies the IBM® Storage Protect node name.
 - **SP Server:** The column specifies the IBM® Storage Protect server name and status.
 - **Monitor Service:** The column allows administrators to start, stop, or restart the monitor service and displays the status of the monitor service.
 - **Recall Service:** The column allows administrators to start, stop, or restart the recall service and displays the status of the recall service.

- **Threshold Migration:** Click the **Threshold Migration** tab to view threshold migration status on different client nodes and volumes. On the **Threshold Migration** page, information is displayed in the tabular format with the following columns:
 - **File Server:** The column specifies the name of a file server.
 - **Volume:** The column specifies the name of a volume on a file server.
 - **HSM Threshold:** The column specifies whether HSM threshold is configured and enabled for a volume.
 - **Low Threshold:** The column specifies the low threshold value for migration.
 - **High Threshold:** The column specifies the high threshold value for migration.
 - **Filespace:** The column specifies the name of the file space.
 - **Expand Icon:** The column allows administrators to update migration configuration by expanding a row. Click the expand icon to view the **HSM Threshold Configuration** dialog box. The dialog box allows administrators to enable or disable migration HSM Threshold Service, configure values for Low Threshold, High Threshold, Minimum File Age, and Minimum File Size, select Filespace and Management Class, and configure other migration options such as Backup, Access Control Lists, and Alternate Data Streams. To create a new file space, click **New File Space** button in the **HSM Threshold Configuration** dialog box.
- **Migration Jobs:** Click the **Migration Jobs** tab to view the status of migration jobs on different client nodes and volumes. The **Create Job** and **Delete Job** options are available on the **Migration Jobs** page to allow administrators to create and delete migration jobs. You can also update migration job configuration by clicking the expand icon available for each volume. On the **Migration Jobs** page, information is displayed in a tabular format with the following columns:
 - **File Server:** The column specifies the name of a file server.
 - **Job:** The column specifies the name of a migration job.
 - **Filespace:** The column specifies the file space name for which a migration job is configured.
 - **Backup:** The column specifies whether the backup is enabled for a migration job.
 - **ACL:** The column specifies whether the file security attributes are migrated when a file is migrated.
 - **ADS:** The column specifies whether Windows alternate data streams data is migrated when a file is migrated.
 - **Location:** The column specifies the migration candidate location.
 - **Expand Icon:** The option allows administrators to update the migration job configuration. Click the expand icon to view the **HSM Migration Job Configuration** dialog box. The dialog box allows administrators specify Job Name, Migration Candidate File Scanner Root Directory, Migration Candidate File Scanner, option to exclude subdirectory in migration, apply file filters and actions for migrated files, select Filespace and Management Class, and configure additional migration options such as Backup, Access Control Lists, and Alternate Data Streams.
- **Volume Analysis:** Click the **Volume Analysis** tab to view the status of scan jobs on different client nodes and volumes. You can also update scan job configuration by clicking the edit icon available for each volume. On the Volume Scan Jobs page, information is displayed in a tabular format with the following columns:
 - **File Server:** The column specifies the name of a file server.
 - **Volume:** The column specifies the name of a volume.
 - **HSM Analysis:** The column specifies whether HSM analysis is configured and scheduled for a volume.
 - **Scan Interval:** The column specifies the time interval between two consecutive volume scans.
 - **Scan Start:** The column specifies the scan start time.
 - **Next Scan:** The column specifies the next scan start date and time.
 - **Expand Icon:** The option allows administrators to update volume analysis configuration. Click the expand icon to view the **HSM Volume Scan Configuration** dialog box. The dialog box allows administrators enable or disable **HSM Volume Scan** for the selected volume and specify values for **Next Scan** and **Scan Interval**.

Support page

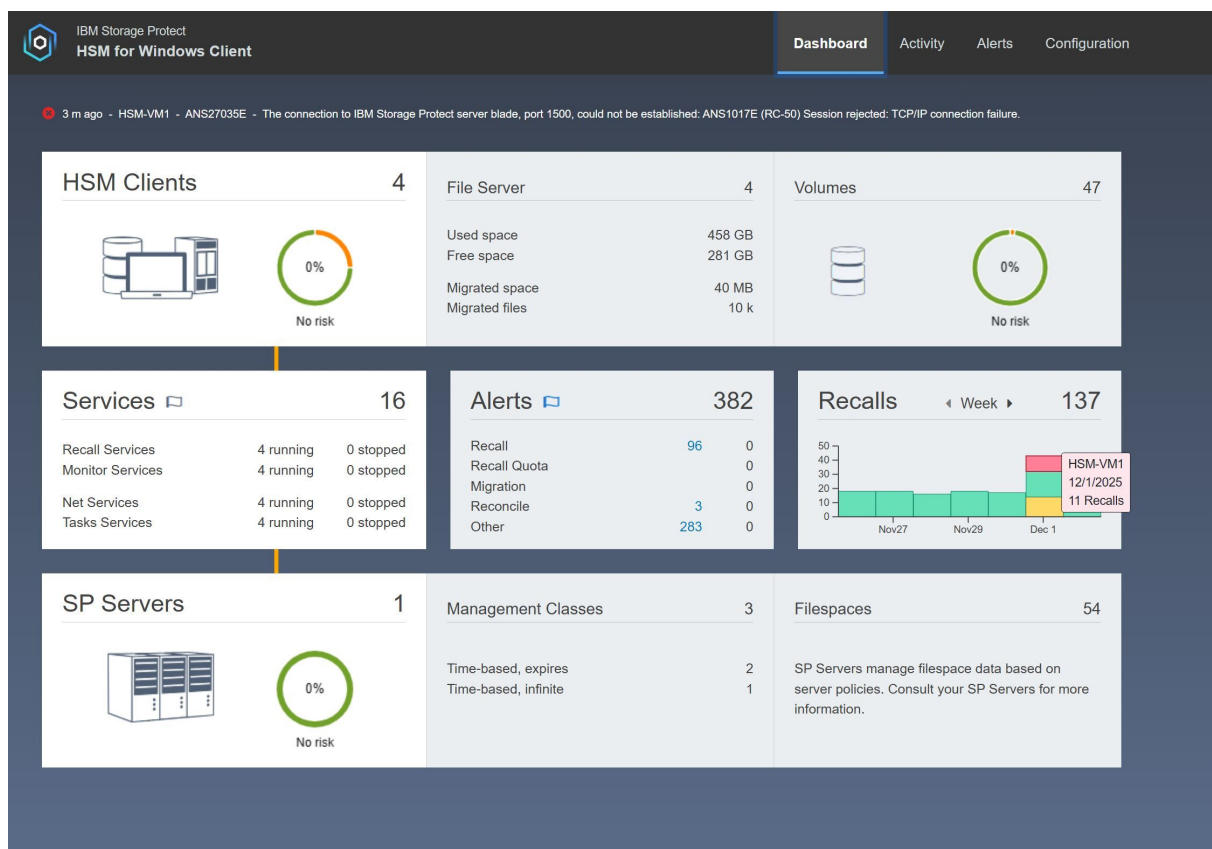
The **Support** page shows an overview of all managed client nodes. The page allows administrators to configure application trace and log settings for different services. You need to select a client node and configure settings for a service to generate and download a package of log files.

HSM for Windows dashboard overview

The **IBM HSM for Windows Client** dashboard provides a consolidated, real-time overview of the operational status of all connected HSM clients and server components. It is designed as a quick-assessment page that highlights system health, ongoing activities, service states, capacity usage, alerts, and the overall condition of the HSM infrastructure.

The **IBM HSM for Windows Client** dashboard displays various tiles, each focusing on a specific aspect of the environment such as storage use, service activity, alerts, and the status of connected Storage Protect (SP) servers. Each tile provides a summary metrics and the ability to drill-down entry points for more detailed analysis in other sections of the user interface.

At the top of the dashboard, global notifications, or error messages such as connection failures are displayed, to help administrators quickly identify systemic issues that may require attention.



Tiles description

- **HSM Clients**: The **HSM Clients** tile displays the number of registered and connected HSM clients. A visual health indicator displayed as green, orange or red with percentage, indicates the current risk level.
- **File Server**: The **File Server** tile displays capacity utilization of the connected file servers used by the HSM environment as following:
 - Used space
 - Free space
 - Migrated space
 - Migrated files

This allows administrators to evaluate storage consumption at a glance.

- **Volumes:** The **Volumes** tile displays the number of available HSM-managed volumes. A risk indicator displayed as green, orange or red with percentage, highlights potential volume-related issues.
- **Services :** The **Services** tile displays the operational status of core HSM services as following:
 - Recall Services
 - Monitor Services
 - Net Services
 - Tasks Services

For each service type, the tile displays the number of instances running or stopped. Thus, enabling the administrators identify service disruptions immediately.

- **Alerts:** The **Alerts** tile provides an overview of active alerts, grouped by the category. A count is shown for each of the following categories:
 - Recall
 - Recall Quota
 - Migration
 - Reconcile
 - Other
- **Recalls:** The **Recalls** tile displays the recent recall activities on the weekly, daily, or hourly biases. The display helps to understand file recovery activity patterns and detect unusual spikes.
- **SP servers:** The **SP servers** tile displays the number and health of configured IBM Storage Protect (SP) servers. A risk indicator displayed as green, orange or red with percentage summarizes the overall server status.
- **Management Classes:** The **Management Classes** tile lists the management class types defined for client operations such as:
 - Time-based, expires
 - Time-based, infinite

The tile displays number of configured classes of each type.

- **File spaces:** The **File spaces** tile displays the number of file spaces managed by the SP servers. A short description explains that SP Servers manage files pace data based on server policies.

Purpose and usage

The **IBM HSM for Windows Client** dashboard provides an immediate, concise overview of system health and activity. Administrators can primarily use it for the following:

- Assess risk and detect operational issues
- Monitor storage usage and volume health
- Track service status and troubleshoot outages
- Review alert categories and address problems
- Observe client recall activity trends
- Validate the availability of SP servers and file spaces

Installation of the HSM for Windows™ client

The backup-archive client uses the IBM® Storage Protect API, which is installed when you install the backup-archive client. Install, configure, and register this client before you install, configure, and use the HSM for Windows™ client.

Planning to install the HSM for Windows™ client

Plan the necessary hardware and software, and consider compatibility with other software.

Hardware and software requirements

HSM for Windows™ client has hardware and software requirements.

For current hardware and software and requirements, see [Technote 1319299](#).

National language environments

When you install or uninstall the HSM for Windows™ client, all languages are installed or uninstalled at the same time. You cannot install the HSM for Windows™ client to a path that contains national language characters.

Compatibility with other software

There are restrictions with file names length and cluster support.

File name limitations

The length of file names is limited by IBM® Storage Protect API, and by Windows® Explorer when using the HSM for Windows™ client GUI.

The length of a file name that is migrated by the HSM for Windows™ client cannot exceed 256 bytes. The path length (the API high-level qualifier) cannot exceed 1024 bytes. A path and file name includes the file server name, the volume, and the directory portion of the full Uniform Naming Convention (UNC) name, for example `\FILESERVER\E:\directory\filename.ext`. The Unicode representation of a character can occupy several bytes, so the maximum number of characters that a file name might contain can vary.

When using the HSM for Windows™ client GUI, path names can be a maximum of 254 characters only. For path names that exceed 254 characters, you must use the `dsmlc.exe` command from a Command Prompt window.

Cluster environment limitations

There are configuration limitations for the HSM for Windows™ client in a cluster environment. Migration jobs must be started again manually when a cluster node fails.

The HSM for Windows™ client supports the following cluster environments:

- A Microsoft® cluster (MSCS) environment with the following configurations:
 - Local volumes that are mounted into local volumes
 - Cluster volumes that are mounted into cluster volumes

Note: In these configurations, both cluster volumes belong to the same cluster resource to guarantee that both are always online on the same cluster node.

You cannot use the following configurations because in these configurations HSM for Windows™ client cannot recall migrated files after failover:

- Cluster volumes that are mounted into local volumes
- Local volumes that are mounted into cluster volumes

When a migration job is running and the cluster node fails, the job is interrupted. You must restart the migration job manually on the next node. When you start the job on the next node, the job continues from the point when the node failed.

Tip: Create a similar job or migration list on the next node in the cluster.

Extended attributes limitations

Extended attributes are not migrated.

Due to a restriction of the NTFS file systems, extended attributes and reparse points are mutually exclusive. Because the HSM for Windows™ client uses reparse points, files with extended attributes cannot be migrated.

Antivirus software

Although HSM for Windows® is tested with popular anti-virus programs, there are several caveats.

Note:

- Be sure a virus scan runs on files before they are migrated.
- Updates of virus signatures and antivirus scan engines can lead to different behavior with the HSM for Windows™ client. During any troubleshooting, always ask the question "What changed?" and take special consideration of antivirus updates.
- Use antivirus software that supports sparse or offline files. Be sure it has a setting that allows it to skip offline or sparse files to avoid unnecessary recall of migrated files.
- The HSM for Windows™ client has been successfully tested for compatibility with the following programs with the specified settings:
 - McAfee VirusScan Enterprise 7.0 and 8.0
 - Symantec AntiVirus 8.0 and 9.0 Corporate Edition with the following setting:
 - Under **Scan Advanced Options > Storage migration options**, check **Skip offline and sparse files**.
 - Symantec AntiVirus 10.0 Corporate Edition with the following two settings:
 - Under **Scan Advanced Options > Storage migration options**, check **Skip offline files**.
 - Under **Autoprotect Advanced Options > Scan files when**, clear **Opened for backup**.

Restrictions for rolling back

Files that are migrated with one version of the HSM for Windows™ client might not be compatible with other versions of the HSM for Windows™ client.

After IBM® Storage Protect HSM for Windows™ version 8.1 is installed and in use, do not roll back to an earlier version or to a fix-pack earlier than version 7.1.6.

Stub files that are created with version 8.1 and higher are not compatible with versions earlier than 7.1.6. versions earlier than fix-pack 7.1.6 of the HSM for Windows™ client cannot recall files from stubs that are created with version 8.1 or version 7.1.6 and later versions.

Preparing for the installation

You can prepare for installation by distributing the installer to the network.

Installing the HSM for Windows™ client

You can install the HSM for Windows™ client from the product installation media.

Before you begin

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

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

Procedure

1. Download the appropriate package file from one of the following websites.
 - Download the HSM for Windows™ client package from [Passport Advantage](#) or [Fix Central](#).
 - For the latest information, updates, and maintenance fixes, go to the [IBM Support Portal](#).
2. Install the product by using the compressed installation file that you download from Passport Advantage®.
 - a. Copy the downloaded compressed installation package to a local disk or to a network-accessible share. Be sure to extract the installation files to an empty directory.
 - b. To extract the installation files, double-click the compressed installation package.
 - c. By default, the uncompressed files are stored in the displayed folder \DISK1. If the installation program detects files from another client installation attempt in this directory, you are prompted about whether to overwrite the old files. If you receive this prompt, enter A to overwrite the existing files; this selection ensures that only the files from the current installation are used.
 - d. Double-click the `spinstall.exe` file to start the client installation program.
 - e. Select a language to use for this installation and click **OK**.

What to do next

After installing the HSM for Windows™ client, you must configure your connection with the IBM® Storage Protect server before you can use HSM for Windows™ client.

Installing and configuring the HSM for Windows™ client in a cluster environment

You can install the HSM for Windows™ client on a Microsoft cluster server (MSCS). With appropriate configuration, the HSM for Windows™ client can manage migration during failover and failback processing.

Before you begin

You must install, configure, and register the IBM® Storage Protect backup-archive client before you configure and use the HSM for Windows™ client.

You must register the cluster node name with the IBM® Storage Protect server as a client node.

About this task

Start the HSM for Windows™ client graphical user interface (GUI) by entering the `dsmgui.exe` command in the HSM for Windows™ client installation directory.

Procedure

1. Install the IBM® Storage Protect backup-archive client and the HSM for Windows™ client on a local drive on each cluster node.
2. Specify the **clusternode** option in the `dsm.opt` client options file of each backup-archive client. Example code in a `dsm.opt` client options file:

```
...
TCPPOPT 1500
PASSWORDACCESS GENERATE
NODENAME WS2008R2CLUSTER
CLUSTERNODE YES
...
```

3. Run the HSM for Windows™ client configuration wizard to configure an HSM client.
 - a. When the **Cluster Configuration** window opens, copy the IBM Storage Protect™ server administrative command `grant proxynode target=targetname agent=agentname` from the **Cluster Configuration** window.
 - b. Run the command on the IBM Storage Protect™ server console, operations center command prompt, or administrative client command prompt.
4. Repeat step 3 on each cluster node.

HSM in cluster environments

Installation, configuration, and use of the HSM for Windows™ client in a cluster environment requires special considerations.

The HSM for Windows™ client manages threshold migration and reconciliation during failover and fallback.

A migration job that is running when a node fails must be manually started on another node of the cluster. When the job is started on another node, files that were processed before the node failed are not processed again. Only the remaining files are migrated. You can configure the same or similar jobs on several nodes of a cluster.

The HSM for Windows™ client must be installed on each cluster node on which files are migrated and recalled. For example, assume that you have a three node cluster. You plan to migrate data from one cluster volume. If this cluster volume is available on only node1 and node2, you need to install the HSM for Windows™ client on only node1 and node2. If the volume can fail over to node3, you must install the HSM for Windows™ client on node3, too.

Each HSM for Windows™ client uses its own node name to authenticate with the IBM® Storage Protect server. By default, the IBM® Storage Protect node name for the computer is the computer host name. But you can change that name when you run the initial configuration wizard.

To access the data from the cluster volumes on all nodes, the data is stored on the IBM® Storage Protect server under a common node name. This common node name must be the cluster name. You must grant access for each node to the common cluster node name by using the **grant proxynode** command. The configuration wizard shows you the appropriate command to be run on the IBM® Storage Protect server.

Each HSM for Windows™ client has its own set of configuration data. The configuration data and the migration jobs and log files are stored by default in subdirectories of the installation directory. You can configure the directory that contains the jobs file to a common directory that is accessible by other HSM for Windows™ client nodes in the cluster. Do not configure a common directory for the logs and list files or for the configuration file directory and temporary files directory.

The HSM for Windows™ client must be installed on each cluster node to a local drive, like the system drive. The HSM executable files must be available at any time. Do not install the HSM for Windows™ client on a cluster drive.

If you want to use the IBM® Storage Protect backup-archive client, it must be installed, configured, and registered appropriately for an MSCS cluster environment. If you want files to be backed up before migration, the options file must specify the **clusternode=yes** option. For example, assume that your cluster volume is E and your backup-archive client scheduler is configured to run the daily backup with the option file `E:\TIVOLI-TSM\dsm_cluster_E.opt`. Select `E:\TIVOLI-TSM\dsm_cluster_E.opt` as the options file for the backup before migration.

Important: HSM for Windows™ client stores the cluster name as file recall information in stub files. If you change the cluster name, you must apply the appropriate hardware volume mappings before you continue.

If you remove a volume from a cluster and reconfigure it as a local volume on one node, you must use hardware volume mappings to link the local volume to the old cluster volume name.

When you install the HSM for Windows™ client on a cluster system, the HSM services require the cluster services. If the cluster services are not running, the HSM services do not start. After you restart the system, the HSM services attempt to start automatically two times. If the cluster services are not running at the second automatic attempt, you must start the HSM services manually.

If you change a cluster name, only the HSM for Windows™ client GUI starts. Use the GUI to map the new cluster name to the old cluster name. When you confirm the mapping that is created by using the wizard, the HSM for Windows™ client creates hardware mappings from the new cluster name to the old cluster name. The mappings are replicated over the IBM® Storage Protect server to other cluster nodes where the HSM for Windows™ client is installed.

Upgrading the HSM for Windows™ client

Upgrading from an earlier version might require that you complete an upgrade task.

Migrating Windows™ alternate data stream data for files that were migrated before version 7.1.2

To use the IBM® Storage Protect HSM for Windows™ alternate data streams (ADS) feature on all migrated files, you must migrate all files with the version 7.1.2 or later client. Files that contain ADS data and were migrated with an earlier version client must be migrated again with the version 7.1.2 or later client.

About this task

With the HSM for Windows™ ADS feature, you can migrate and retrieve Windows™ ADS data. ADS data in a stub is backed up and can be restored by the backup-archive client.

If you migrated files that contain ADS data with an earlier version client, the migration copy does not contain ADS data. If the stub was backed up, the stub does not contain ADS data. ADS data is not fully protected in these files. If you want full ADS protection, you must migrate the files that contain ADS data again.

Procedure

1. Plan enough space resources and time to rerun all migration jobs.
2. When HSM for Windows™ 7.1.2 or later is installed, rerun all migration jobs.

Result

For the stub files that are migrated with HSM for Windows™ 7.1.2 or later, the migration copies in IBM Storage Protect™ storage contain ADS data. The ADS data that is in stub files is backed up at the next scheduled incremental backup or image backup. After the incremental or image backup, ADS data is fully protected.

Configuring the HSM for Windows® client

This topic indicates when and how to configure the HSM for Windows™ client.

After installing the HSM for Windows™ client, you must configure your connection with the IBM® Storage Protect server before you can use HSM for Windows™ client. The first time you start the GUI, the configuration wizard guides your choices. After the initial configuring the connection to the IBM® Storage Protect server, you can use the configuration wizard at any time to change the initial settings.

The HSM for Windows™ client is installed with default values for regional settings, file recall settings, and the location of configuration, log, and job files. You can change these values at any time with the **Preferences** window.

You can configure migration jobs, threshold migration, or reconciliation at any time after configuring connection with the IBM® Storage Protect server.

After adding new hard disks or volumes to a computer that is already running the HSM for Windows™ client, you must restart the recall service (`hsm.service.exe`) and the monitor service (`hsmmonitor.exe`).

Configuring the connection between the HSM for Windows™ client and the IBM® Storage Protect server

You must configure the connection between the HSM for Windows™ client and the IBM® Storage Protect server before you can use the HSM for Windows™ client.

Before you begin

You must install, configure, and register the IBM® Storage Protect backup-archive client before you configure and use the HSM for Windows™ client.

If the computer is a cluster node, you must complete the following tasks:

- Register the cluster node name with the IBM® Storage Protect server as a client node.
- Add `clusternode yes` to the backup-archive client `dsm.opt` options file.
- For installation and configuration requirements in a cluster environment, see [“Installing and configuring the HSM for Windows client in a cluster environment” on page 25](#)

About this task

The first time that you run the HSM for Windows™ client graphical user interface (GUI), the Configuration wizard displays. The Configuration wizard guides you through the steps to configure a connection between the HSM for Windows™ client and the IBM® Storage Protect server. You can also run the Configuration wizard any time from the **Tools** menu.

Start the HSM for Windows™ client graphical user interface (GUI) by issuing the `dsmsgui.exe` command in the HSM for Windows™ client installation directory.

Procedure

1. In the **Option File Task** page, choose whether to create an options file or update an existing options file. If there is no options file, you must create an options file. Click **Next**.
The HSM for Windows™ client stores configuration information in the `dsm.opt` file that is in the HSM for Windows™ client installation directory. It does not use the `dsm.opt` file that is used by the IBM® Storage Protect backup-archive client.

Attention: Use only the HSM for Windows™ client GUI to change HSM for Windows™ client options. Editing the HSM for Windows™ client `dsm.opt` file by another method risks corrupting the file, and can lead to loss of data.

Password and names of file spaces are also stored and managed separately from the backup-archive client. They are stored and managed with the Windows™ registry entries of the HSM for Windows™ client.

2. In the **TPC/IP Parameters** window, enter the server address and port for the IBM® Storage Protect server. Select TCP/IP options and select **Next**.
Each HSM for Windows™ client can connect to only one IBM Storage Protect™ server for migration. This server can be different from the one that is used by the backup-archive client. If you do not select the box for TCP/IP V4 and TCP/IP V6, the HSM for Windows™ client uses only TCP/IP V4.
3. In the **IBM Storage Protect Authentication** window, enter the IBM Storage Protect™ client node name and click **Next**.
The node name must be registered with the IBM® Storage Protect server. If you want to clearly identify the HSM node as distinct from the backup-archive node, choose a different node name for the HSM for Windows™ client. If the computer is a cluster node, the client node name must be the cluster node name. If you want to keep the logon parameters of the HSM for Windows™ client separate from the logon parameters of the backup-archive client, register the HSM for Windows™ client under a node name different from the node name that is used by the backup-archive client.
4. Configure the HSM for Windows™ client for a cluster configuration.
If the computer is a cluster node, the **Cluster Configuration** window is displayed.
 - a. Copy the IBM Storage Protect™ server administrative command `grant proxynode target=targetname agent=agentname` from the **Cluster Configuration** window.
 - b. Run the command on the IBM Storage Protect™ server console, operations center command line, or administrative client command line.
The value of the target parameter (the cluster node name) and the value of the agent parameter (the client node name) must be registered with the IBM Storage Protect™ server as client nodes.
 - c. Click **Next**.
5. In the **Set or Change Password** window, type the password for the node and click **Next**.
The password was created when the node was registered with the IBM Storage Protect™ server. You can change the password in this panel.
6. In the **IBM Storage Protect Server Connection** window, verify the values that you configured in the previous windows. Click **Apply**.
7. In the **IBM Storage Protect Server Management Class** window, select the management class that is the default when you create a migration job or configuration threshold migration and click **Next**.
The choice is the default for migration jobs and threshold migration, but you can override the default on each operation. Information in the window indicates the suitability of the management class for archived migration copies.
8. In the default file-security window, specify whether file security attributes (ACL) are migrated and retrieved and click **Next**.
The choice is the default for migration jobs, list migrations, selective retrieves, and threshold migrations. You can override the default on each operation in a job file or with a command parameter.
9. In the **Alternate Data Streams** window, specify whether Windows alternate data stream (ADS) data is migrated or retrieved and click **Next**.
The choice is the default for migration jobs, list migrations, selective retrieves, and threshold migrations. You can override the default on each operation in a job file or with a command parameter.
10. In the **Backup Before Migration** window, configure whether files are backed up before they are migrated and click **Next**.
If you choose to back up files before migration, select an options file for the backup. If this option is cleared, the default is to back up before migration. The choice is the default for migration jobs, list migrations, and threshold-migration configurations. You can override the default on each migration job and each threshold-migration configuration.
11. If no file space was registered, the **Initial File Space Registration** window is displayed. Enter the name of the default file space to store migrated files from your client node on the IBM Storage Protect™ server, or select the **Skip file space creation** check box. Click **Next**.
If you enter the name of a file space that does not yet exist, the HSM for Windows™ client creates the file space. Select the **Skip file space creation** check box if you want to create a file space when you define a migration job or threshold migration, or start a list migration.
12. Confirm the settings in the **Completing the IBM Storage Protect HSM Configuration Wizard** window. If all options are correct, click **Finish**. If you must make corrections, click **Back**.

What to do next

When the HSM for Windows™ client connects successfully to the IBM Storage Protect™ server you can configure migration jobs, threshold migration, and reconciliation.
If the computer is a cluster node, you must repeat the configuration for the HSM client on each node of the cluster.

Client password-character restrictions

HSM for Windows™ client passwords are restricted to certain characters. In some cases, passwords are case-sensitive.

Passwords can be up to 63 character in length. Password constraints vary, depending on where the passwords are stored and managed, and depending on the version of the IBM® Storage Protect server that your client connects to.

If your IBM® Storage Protect server is at version 6.3.3 or later, and if you use an LDAP directory server to authenticate passwords

Use any of the following characters to create a password:

```
a b c d e f g h i j k l m n o p q r s t u v w x y z
A B C D E F G H I J K L M N O P Q R S T U V W X Y Z
0 1 2 3 4 5 6 7 8 9 . ! @ # $ % ^ & * _ - + = ` ( )
| { } [ ] : ; < > , ? / ~
```

Passwords are case-sensitive and are subject to more restrictions that can be imposed by LDAP policies.

If your IBM® Storage Protect server is at version 6.3.3 or later, and if you do not use an LDAP directory server to authenticate passwords

Use any of the following characters to create a password:

```
a b c d e f g h i j k l m n o p q r s t u v w x y z
A B C D E F G H I J K L M N O P Q R S T U V W X Y Z
0 1 2 3 4 5 6 7 8 9 . ! @ # $ % ^ & * _ - + = ` ( )
| { } [ ] : ; < > , ? / ~
```

Passwords are stored in the IBM® Storage Protect server database and are not case-sensitive.

If your IBM® Storage Protect server is earlier than version 6.3.3

Use any of the following characters to create a password:

```
a b c d e f g h i j k l m n o p q r s t u v w x y z
A B C D E F G H I J K L M N O P Q R S T U V W X Y Z
0 1 2 3 4 5 6 7 8 9
_ - & + .
```

Passwords are stored in the IBM® Storage Protect server database and are not case-sensitive.

Remember:

On the command line, enclose all parameters that contain one or more special characters in quotation marks. Without quotation marks, the special characters can be interpreted as shell escape characters, file redirection characters, or other characters that have significance to the operating system.

On Windows™ systems:

Enclose the command parameters in quotation marks (").

Command line example:

```
dsmc set password "t67@#$$^&" "pass2><w0rd"
```

Quotation marks are not required when you type a password with special characters in an options file.

Configuring the HSM client to connect to a secondary IBM® Storage Protect server

If the primary IBM® Storage Protect server for the HSM for Windows™ client is unavailable, you can manually configure the HSM for Windows™ client to connect to a secondary server. You can recall files from the secondary IBM® Storage Protect server but cannot migrate files to the secondary server.

Before you begin

The primary IBM® Storage Protect server for the HSM for Windows™ client must be one that replicates client node data.

About this task

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

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

The HSM for Windows™ client, however, does not automatically fail over to the secondary server. You must manually configure the `dsm.opt` file to connect to the secondary server. Any secondary server information in the `replservername` stanza, the `myreplicationserver` option, and the `myprimaryserver` option is ignored by the HSM for Windows™ client.

You can complete some tasks when connected to the secondary server:

- You can recall and retrieve migrated files from the secondary server by using the HSM for Windows™ client.
- You can restore a stub file by using the backup-archive client.
- You must not reconcile the file system with the secondary server.
- You must not migrate files to the secondary server.

Procedure

To configure the HSM for Windows™ client to connect to the secondary server complete the following step:

1. Edit the `dsm.opt` file to specify information about the secondary server.
The following stanza is an example of a secondary server stanza:

| | |
|------------------|--------------------------|
| COMMMethod | TCPip |
| TCPPort | 1500 |
| TCPServeraddress | lifeboat.almaden.ibm.com |
| Passwordaccess | generate |

What to do next

After you complete this steps, restart the HSM for Windows™ client.

You can complete some tasks when connected to the secondary IBM® Storage Protect server:

- You can recall and retrieve migrated files from the secondary server by using the HSM client.
- You can restore a stub file using the backup-archive client.
- You must not reconcile the file system with the secondary server.
- You must not migrate files to the secondary server.

Connect to the primary IBM® Storage Protect server as soon as it becomes available.

Configuring the retention period of migration copies

You can control the period for which migration copies are stored in IBM® Storage Protect storage. If you accept the installed-default data management policy, migration copies can be deleted from IBM Storage Protect™ storage in one year.

Files that are migrated by HSM for Windows™ client are stored as migration copies on an IBM® Storage Protect server. The migration copies are stored in the storage pool that is defined by the archive copy group of the

assigned management class. When migration copies are created in the HSM pool, they are bound to a management class. The migration copies are retained according to the policy specified in the archive copy group of the management class. If the retention period is too short, IBM Storage Protect™ can delete the migration copies on the IBM Storage Protect™ server and leave orphan stubs on the file system. In this case, the migrated files cannot be recalled, and must be restored from backup copies.

If you do not specify a management class for your migration copies, they are bound to the default management class. The default policy values in the archive copy group of the standard management class retain migration copies for only one year.

If the default management class has no archive copy group, the migration copies are retained according to the **ARCHREtention** value defined for the domain.

The archive copy group specifies three attributes that determine the period that migration copies can be retained on the IBM® Storage Protect server.

- **RETVer** determines the number of days to retain a migration copy.
- **REtInit** determines when the **RETVer** attribute is applied.
If **REtInit=Event**, the **RETVer** attribute applies when a HSM for Windows™ client reconciliation process determines that a migration copy is no longer needed. Migration copies are retained like this:
 - a. A stub is deleted from the file system.
 - b. Reconciliation determines that the migration copy on the IBM® Storage Protect server is no longer needed. Reconciliation sends an event notice to the IBM® Storage Protect server.
 - c. When the IBM® Storage Protect server receives the event notice from the reconciliation process, the retention period specified by **RETVer** begins.
 - d. When the retention period specified by **RETVer** ends, the IBM® Storage Protect server marks the file for deletion.
 - e. When the IBM® Storage Protect server runs an expiration process, the migration copy is deleted from the IBM® Storage Protect server.

If **REtInit=CREAtion**, the **RETVer** attribute applies when a migration copy is created. If the **RETVer** period expires before a stub is deleted, IBM® Storage Protect server deletes the migration copy. This leaves an orphan stub on the file system. If a stub is deleted before the **RETVer** period expires, a migration copy is retained like this:

- a. A stub is deleted from the file system.
- b. Reconciliation determines that the migration copy on the IBM® Storage Protect server is no longer needed. Reconciliation sends a deletion notice to the IBM® Storage Protect server.
- c. When the IBM® Storage Protect server receives the deletion notice from the reconciliation process, the IBM® Storage Protect server immediately marks the migration copy for deletion.
- d. When the IBM® Storage Protect server runs an expiration process, the migration copy is deleted from the IBM® Storage Protect server.

After a copy group is defined, the **REtInit** value cannot be updated.

- **REtMin** determines the minimum period to retain a migration copy after it is created. This attribute applies only when **RETVer=Event**.

Choose a management class with an archive copy group that meets your data retention needs.

When you configure the connection between the HSM for Windows™ client and the IBM® Storage Protect server, you can specify a management class. This management class becomes the default management class for new migration jobs and new threshold migration configurations. You can specify a different management class for migration when you configure a job or threshold migration, and when you start a migration using `dsmc1c.exe`. The management class that you specify when you configure a job or threshold migration overrides the default management class for migration. The management class that you specify when you start a migration using `dsmc1c.exe` overrides the configured management class for migration.

Jobs and threshold migration that were configured prior to version 6.1.3 did not specify a management class, and they used the default management class for the policy set. Those jobs and threshold migration continue to use the default management class for the policy set until you reconfigure them. Note that the default management class for the policy set can be the same as the default management class for new migration jobs and threshold configuration, but is not necessarily the same.

Changing the retention period of migration copies

You can change the retention period of migration copies that are stored on an IBM® Storage Protect server.

When files are migrated by HSM for Windows™ client, they are bound to a management class. The retention period of migration copies are determined by the archive copy group settings of that management class. To change the retention period of the migration copies, you must change the archive copy group settings.

There are several ways you can change archive copy group settings. The simplest change is to update the archive copy group settings of the management class that is currently bound to the migration copies. Although the change is simple, the change affects all archive copies that are bound to this management class. This can include copies of files that are archived by the backup-archive client. And you are limited because when you update an archive copy group, you cannot change the **RETInit** value.

A more complex change involves creating a new domain for HSM for Windows™ client migration copies. IBM Storage Protect™ policy allows many ways to change the archive copy group settings, and you can choose the option that works best for your business. The following recommendations assume that migration copies are currently bound to the default management class. This assumption would be true for migration copies that are created by HSM for Windows™ Version 6.1.2 and earlier. These suggestions can be modified to account for migration copies that are not currently bound to the default management class.

Define a new policy domain that isolates the HSM for Windows™ client from other client nodes.

Define a new policy domain just for the HSM for Windows™ client. Define a policy set for the new domain. Define a new management class with an archive copy group that specifies an appropriate retention period for migration copies. Assign the new management class as the default for the new policy domain and policy set. Validate and activate the policy set. Update the HSM for Windows™ client node to become a member of the new policy domain.

As a result, all migration copies on the IBM Storage Protect™ server that are associated with the HSM for Windows™ client node and that were previously bound to the old default management class are rebound to the new default management class.

If the HSM for Windows™ client node name is the same as the backup-archive client node name, this change can also affect the archive copies created by the backup-archive client.

This solution works for all versions.

Define a new default management class for the existing domain

Define a new management class with an archive copy group that specifies an appropriate retention period for migration copies. Assign the new management class as the default for the existing policy domain and policy set.

As a result, all migration copies on the IBM Storage Protect™ server that are associated with the existing policy domain and that were previously bound to the old default management class are rebound to the new default management class. This change can affect the migration copies of all nodes that are members of the policy domain.

This solution also works for files that were migrated with HSM for Windows™ client versions earlier than 6.1.3. Files that were migrated with such earlier HSM for Windows™ client versions are bound to the default management class.

Recall and remigrate files with a new management class

Define a new management class with an archive copy group that specifies an appropriate retention period for migration copies. The new management class does not have to be the default for the active policy set. Recall all migrated files. Delete the existing file spaces. Migrate the files again, and specify the new management class.

As a result, the migration copies on the IBM Storage Protect™ server that were created by the HSM for Windows™ client are bound to the new management class. This change does not affect the archive copies that were created by the backup-archive client. This process can cause significant network traffic and use significant local storage resources.

Configuring a new file space

You can create new file spaces on the IBM® Storage Protect server directly from the HSM for Windows™ client GUI.

About this task

Use the steps in this task to create a new file space:

Procedure

1. To create a new file space select **Tools > Create New File Space**.
2. Enter a name for the new file space.
3. Select the **OK** button.

Configuring regional settings

Use the **Regional Settings** tab of the **Preferences** window to set your language, time format, date format, number format, and define if you want log, list, and trace files in Unicode.

Before you begin

Note: You must restart the HSM for Windows™ client GUI for any changes to become effective.

Procedure

1. Select **Tools > Preferences** and then select the Regional Settings tab.
2. Make changes as needed and select the **OK** button.

Excluding Windows™ alternate data stream names

You can exclude Windows™ alternate data stream (ADS) data by name. ADS names in the exclude list are excluded from HSM operations.

About this task

To create a list of Windows™ alternate data stream names that are excluded from HSM operations, complete the following steps.

Procedure

1. In the HSM for Windows™ client GUI, click **Tools > ADS Exclude List**.
The **ADS Exclude List** window is displayed.
2. Click **Create**, type an ADS name, and click **OK**.

HSM advanced parameters and preferences settings

Although most parameter default settings are appropriate, you can customize some settings.

“Advanced parameters descriptions” on page 36 displays the advanced parameters. For all parameters except the Timeout parameter, the Parameter name column shows you the parameter name and Windows® registry path from the end of this common path: HKLM\SOFTWARE\IBM\ADSM\CurrentVersion\HsmClient\. The Timeout parameter full path is listed in the Parameter name column.

| Table 3: Advanced parameters descriptions | | | |
|--|--|-----------------------|---|
| Parameter name | description | Default | Notes |
| HKLM\SYSTEM\CurrentControlSet\Services\ithsmdrv\Parameters\Timeout | The File System Filter Driver returns an error when this time elapses and a recall process has not yet started. If the recall process starts within this time, no error is returned. The start time is when the recall thread picks up the recall order. Time waiting for a device or reading data is not considered. The end time of the recall process is not considered. The time is measured in seconds. | 300 | The error is returned when the recall service is too busy, and the recall quota has not been reached. This can occur when many recall processes are running at the same time. |
| dsmclic\FileAttributesFilter | Configures the registry to prevent files with certain attributes from migration. Affects the dsmclic.exe command. | 6 - hidden and system | Change this parameter only on technical advice from IBM®. |
| dsmgui\FileAttributesFilter | Configures the registry to prevent files with certain attributes from migration. Affects the dsmgui.exe command. | 6 - hidden and system | Change this parameter only on technical advice from IBM®. |
| hsmmonitor\FileAttributesFilter | Configures the registry to prevent files with certain attributes from migration. Affects the hsmmonitor.exe command. | 6 - hidden and system | Change this parameter only on technical advice from IBM®. |
| dsmclic\DirectoryAttributesFilter | Configures directories with certain attributes that are generally not entered for selecting files for migration. Affects the dsmclic.exe command. | 6 - hidden and system | Change this parameter only on technical advice from IBM®. |
| dsmgui\DirectoryAttributesFilter | Configures directories with certain attributes that are generally not entered for selecting files for migration. Affects the dsmgui.exe command. | 6 - hidden and system | Change this parameter only on technical advice from IBM®. |
| hsmmonitor\DirectoryAttributesFilter | Configures directories with certain attributes that are generally not entered for selecting files for migration. Affects the hsmmonitor.exe command. | 6 - hidden and system | Change this parameter only on technical advice from IBM®. |

File location preferences

Use the HSM for Windows™ client GUI **Preferences** window **Path Configuration** tab to define file locations.

Access the **Preferences** window **Path Configuration** tab by selecting HSM for Windows™ client GUI. Select **Tools > Preferences > Path Configuration**.

The **Path Configuration** tab contains fields that indicate the location of the following files:

- Configuration files
- Migration job files

- Move job files
- Temporary files

Move Settings

You can configure the bandwidth that is used for moving stub files. You can also configure how many stub files are identified before a move process begins.

Use the HSM for Windows™ client GUI **Preferences** window **Move Settings** tab to configure two move settings:

Bandwidth

The **Bandwidth** value controls what percent of time the HSM for Windows™ client spends on move operations. For example, if you set **Bandwidth**=40% and a move operation takes 20 milliseconds, the HSM for Windows™ client pauses for 30 milliseconds before the next move operation starts. The total elapsed time is 50 milliseconds, the move operation is 20 milliseconds (40%) of the elapsed time.

Stub Files

The **Stub Files** value controls how many stub files are identified before a move operation begins. The HSM for Windows™ client moves the stub files in an optimal order to minimize the number of tape mounts and seeks. When the list of stub files is large, more files can be moved with fewer tape mounts. However, it takes more time for the HSM for Windows™ client to identify a large number of stub files. A larger value improves the efficiency of the move process, but delays the start of the move operation. The value can be 1 to 50,000. The default is 5,000.

File recall quotas

You can define file recall quotas to limit the number of possible file recalls during a time span. You can define a system-wide default quota and define quotas for particular Windows® user and group accounts.

Several quotas can be defined for one user account:

- A user account quota can be defined.
- A user account can be a member of one or more group accounts for which a group account quota is defined.
- A default quota can be defined.

The quota that applies to a user account is the *effective quota*.

User account quotas define the allowed number of file recalls in a time span for an individual user account. If a user account quota is defined, only this quota applies to the user account. Default and group account quotas are overridden by a user account quota.

Group account quotas define the allowed number of file recalls in a time span for each user account in a group. If a user account is a member of two or more groups and has no user account quota, the group with the least restrictive quota applies to the user account.

The default quota applies to user accounts for which no group account quota or user account quota is defined.

You can define quotas for global groups and for universal groups. You cannot define quotas for local domain groups. Local-domain group quotas that were defined in previous versions of HSM for Windows™ are ignored.

Quotas can be updated at any time with the HSM GUI. The update is effective immediately without restarting the HSM for Windows™ client. The update is displayed in the **Live Quotas** tab of the **Recall Quotas** window after the user recalls a file.

The HSM for Windows™ client compares the file recalls quota with the actual file recalls during a time span. The time span is a moving window. For example, assume that you define a quota of five files per 60 seconds. When a user tries to recall a file, the HSM for Windows™ client compares the file recalls quota with the number of file recalls in the previous 60 seconds. If the user recalled five files in the previous 60 seconds, the user cannot recall another file until more time elapses. When less than five files were recalled in the previous 60 seconds, the user can recall another file.

When a user reaches a file recall quota, a subsequent file recall request is rejected. The HSM for Windows™ client returns the code STATUS_FILE_IS_OFFLINE. The behavior of the calling application depends on the response of the calling application to the return code.

When a quota is reached, the end user on the client workstation might not know why the request for access is denied. The HSM for Windows™ client writes a warning message to the hsm service log file and writes a record in the hsm service listing file with result value 'Quota denied'. The HSM for Windows™ client cannot communicate to the end user on a client system that accesses a share on the Windows® server. Administrators can communicate the recall quota and consequences to end users by using a FAQ document, for example.

Quotas affect only the recall of migrated files from users that access stub files. Quotas do not have any influence on retrieving files with the HSM for Windows™ client GUI.

If a user reaches the quota, you can reset the file recall counter. You can reset the file recall counter with the **dsmquota.exe** command or **Live Quotas** tab of the **Recall Quotas** window of the HSM for Windows™ client GUI.

The quota configuration is stored in the HSM for Windows™ client installation directory in `\config\quota.cfg`. After you change quotas, a backup of `quota.cfg` is saved in the backup directory `\config\backup\quota.cfg`.

Viewing and changing the default file-recall quota

Use the HSM for Windows™ client GUI to view and change the default file-recall quota.

Procedure

1. Select **Tools > Recall Quotas** to open the **Recall Quotas** window.
2. Select the **Default Quota** tab.
3. Change the quota.
 - a. Select one of the following options:
 - Unlimited Recalls**
There is no limit for file recalls.
 - No Recalls**
No recalls are allowed.
 - Configure Quota**
You must enter the number of files and the time span.
 - b. Click **OK** to change the default quota.

Viewing and changing a group file-recall quota

Use the HSM for Windows™ client GUI to view and change a group account quota.

About this task

The effective quota for a user account is determined by the hierarchy of quota types and from the quota definitions that apply to the user account. Quota types have the following hierarchy:

- A group account quota overrides the default quota.
- The highest group account quota overrides other group account quotas.
- A user account quota overrides a group account quota.

Procedure

Complete the following steps.

1. Select **Tools > Recall Quotas** to open the **Recall Quotas** window.
2. Select the **Group Quotas** tab.
3. Filter the group accounts by domain and group account name.
 - a. In the **Look in** list, select a domain.

- b. In the **Filter** field, type a group account name pattern. You can use wildcard character * to replace one or more characters and ? to replace one character.
 - c. Click **Find Now** to display group accounts that meet the domain and name criteria.
4. Change a group account quota.
 - a. Select a group account and click **Change Quota**.
The **Recall Quota Editor** window opens.
 - b. Select one of the following choices:

No Group Quota Definition

Do not apply this quota definition. The quota definition is not used to calculate the effective quota for a user account.

Unlimited Recalls

There is no limit for file recalls.

No Recalls

No recalls are allowed.

Configure Quota

You must enter the number of files and the time span.
 - c. Click **OK**.
 - d. Click **OK** to change the quota.

Viewing and changing a user file-recall quota

Use the HSM for Windows™ client GUI to view and change a user account file-recall quota.

About this task

The effective quota for a user account is determined by the hierarchy of quota types and from the quota definitions that apply to the user account. Quota types have the following hierarchy:

- A group account quota overrides the default quota.
- The highest group account quota overrides other group account quotas.
- A user account quota overrides a group account quota.

Procedure

Complete the following steps.

1. Select **Tools > Recall Quotas** to open the **Recall Quotas** window.
2. Select the **User Quotas** tab.
3. Filter the user accounts by domain and user name.
 - a. In the **Look in** list, select a domain.
 - b. In the **Filter** field, type a user account name pattern. You can use wildcard characters asterisk (*) and question mark (?).
 - c. Click **Find Now** to display users that meet the domain and name criteria.
4. Select a user and click **Effective Quota**.
The **Effective User Recall Quota** window shows all quota definitions for the user account and the effective quota for the user account.
5. Change a user account quota.
 - a. Select a user account and click **Change Quota**.
The **Recall Quota Editor** window opens.
 - b. Select one of the following choices:

No User Quota Definition

Do not apply this quota definition. The quota definition is not used to calculate the effective quota for a user account.

Unlimited Recalls

There is no limit for file recalls.

No Recalls

No recalls are allowed.

Configure Quota

You must enter the number of files and the time span.

- c. Click **OK**.
- d. Click **OK** to change the quota.

Viewing and resetting file-recall counters

Use the HSM for Windows™ client GUI to view live file-recall counters. You can reset file recall counters.

Before you begin

The IBM Storage Protect™ HSM Recall Service must be running. If the IBM Storage Protect™ HSM Recall Service is not running, live quota information is not available.

About this task

Live quota information is periodically refreshed. You can change the frequency at which the view is refreshed in **Tools > Preferences > Recall Service**.

Procedure

1. Select **Tools > Recall Quotas** to open the **Recall Quotas** window.
2. Select the **Live Quotas** tab.
3. Filter the user accounts by domain and user name.
 - a. In the **Look in** list, select a domain.
 - b. In the **Filter** field, type a user account name pattern. You can use wildcard characters asterisk (*) and question mark (?).
 - c. Click **Refresh**.

The file recall counter and quota are displayed for user accounts for which all of the following are true:

- The user account is found in the domain.
- The user name matches the filter.
- The file recall quota for the user account is finite and greater than 0.
- The file recall counter for the user account is greater than 0. After you change a quota definition, live-quota information for the user account is not displayed until the user recalls a file.

The live quota information is of the format **1 of 3 recalls**, where 1 is the recall counter and 3 is the recall quota. The recall counter indicates the file recalls that are within the time span that is defined for the quota. The file recall counter changes as the user recall files and as the time-span window changes. The button name changes from **Refresh** to **Pause**.

4. Click **Pause**.
Live updates are paused.
5. Select a user account and click **Reset Counter**.
The file recall counter is reset to 0 for the user account.
6. Select a different domain or a different name filter.
When live updates are refreshed, user accounts that match the domain and name criteria are displayed.

Recall service settings

Use the HSM for Windows™ client GUI to define the recall service settings. You can set the number of concurrent connections to the IBM Storage Protect™ server and the time period for closing connections and deleting obsolete recall quota counters.

Access the **Recall Service** tab of the **Preferences** window by selecting **Tools > Preferences > Recall Service**.

Restriction: Change the value of **Threads** only when required by IBM® Software Support. The **Threads** value determines the maximum number of concurrent connections for the recall service. The default is 4 and the maximum is 64.

You can set the time to close an idle connection to the IBM Storage Protect™ server. The default value is 600 seconds.

Note: If a file is recalled from a tape, the connection is reset to ensure that the tape is not locked after the recall.

You can change the frequency at which the view is refreshed in the **Live Quotas** window. The default value is 2 seconds.

You can set the interval to delete expired quota entries. To determine file recall quotas, the recall service creates a record for every file recall. Periodically, a collection routine runs to eliminate obsolete table entries. Running the collection routine frequently saves computer memory but requires more computer processing. The default value is 60 minutes. The minimum value is 10 minutes, and the maximum value is 10080 minutes.

The frequency of running the collection routine does not affect recall performance. The collection routine is not a performance tuning tool.

Tracing preferences

HSM for Windows™ client processing, from both the GUI and the commands, creates several log files, trace files, and list files.

You can set the logging levels, log file sizes, and log file locations in the **Trace Preferences** window in the HSM for Windows™ client GUI. You can also set the log levels with HSM for Windows™ client commands. You cannot set the log file location or the size with HSM for Windows™ client commands.

In normal production, the defaults log values are sufficient. The default level records warnings and errors and does not record trace-level messages. Increase the logging level only when you must complete advanced diagnostic tasks. The **Severe** and **Error** logging levels are active by default and cannot be deactivated.

When you change log levels in the **hsm service**, **hsm tasks**, **hsm monitor** or **ds m gui** tab, you are not required to restart those programs for those settings to become active. Other changes might require a restart. You are notified when a restart is required.

There are three types of settings you define for the logs: their recording level, their size, the log file location. To access these settings from the HSM for Windows™ client GUI, select **Tools > Trace Preferences**.

[“Tracing preferences: Trace levels” on page 41](#) describes the trace levels settings.

| Table 4: Tracing preferences: Trace levels | |
|--|---|
| Field | Description |
| Severe | Records HSM Windows® messages that are categorized as severe. |
| Error | Records HSM Windows® messages that are categorized as errors. |
| Warning | If checked, records HSM Windows® messages that are categorized as warnings. |
| Info | If checked, records HSM Windows® messages that are categorized as information only. |

| Field | Description |
|----------------|---|
| Trace | If checked, turns on the tracing of program events and is used for advanced diagnostic tasks or for problem analysis. |
| Debug | If checked, records special debugging information and codes is used for advanced diagnostic tasks or for problem analysis. |
| Library | If checked, records specific library information and is used for advanced diagnostic tasks or for problem analysis. |
| Dump | If checked, records more information about issues and is used for advanced diagnostic tasks or for problem analysis. |
| Events | If checked, records diagnostic information such as function entries and exits. |
| Flush | If checked, records each message to disk before processing continues instead of buffering them. This records all messages one-by-one but might affect system performance. Use this setting for advanced diagnostic tasks. |
| Default | Returns the settings in the Trace Levels section of this window to their default values. |
| Full | Returns all available logging and tracing levels. |

“Tracing preferences: Trace file size” on page 42 describes the trace file size settings.

| Table 5: Tracing preferences: Trace file size | |
|---|---|
| Field | Description |
| Wrap the trace file | <p>Defines whether the trace file wraps. By default, the option is set and the trace files wraps when the maximum file size is reached.</p> <p>When the option is cleared, the trace file does not wrap. All trace records are saved:</p> <ul style="list-style-type: none"> The HSM for Windows™ client appends the current date and time to the trace file name. When the trace file reaches the maximum file size, the trace file is saved and a new trace file is created. The HSM for Windows™ client appends the current date and time to the new trace-file name. <p>After you set or clear the option, the new setting takes effect immediately for the dsmsgui application and the hsmmonitor, hsmtasks, and hsmservice services. For commands, the new setting takes effect the next time that you run a command.</p> |
| Maximum file size | Sets a size limit in megabytes for the selected trace file. The default is 10. |
| File wrapping at | Defines the percentage of the log file that is kept when the Maximum file size value is reached. The default is 66. |

“Tracing preferences: Log file size settings” on page 42 describes the log file size settings.

| Table 6: Tracing preferences: Log file size settings | |
|--|--|
| Field | Description |
| Maximum file size | Sets a size limit in megabytes for the selected log file. The default is 10. |
| File wrapping at | Defines the percentage of the log file that is kept when the Maximum file size value is reached. The default is 66. |

The **Path Configuration** tab contains three text boxes where you select the path of the three different files: trace files, log files, and list files. Click **Browse** to select an existing directory.

Managing space with HSM for Windows®

You can manage space on Windows® file servers by creating and running migration jobs, and by configuring threshold migration.

You can manually retrieve migrated files with the HSM for Windows™ client or with the IBM® Storage Protect backup-archive client.

Changes on your file system need to periodically be reconciled with the IBM® Storage Protect server.

Migration jobs

A migration job specifies which files to migrate and whether to leave a stub file on the originating file system.

You can specify the files to migrate by using the HSM for Windows™ client GUI or the HSM for Windows™ client `dsmc1c.exe` command.

With the HSM for Windows™ client GUI, you can browse local file systems. You can exclude or include parts of the directory structure in a migration job. For each selection, filters can be applied to include or exclude files. The inclusion or exclusion can be based on file criteria:

- File type
- File size
- File creation date
- File modification date
- File access date

Each migration job is stored in an XML structured job file. The actual migration can be scheduled by using any standard scheduler, or it can be started manually from a **Command Prompt** window. In addition, the HSM for Windows™ client administrator can start a migration job directly from the HSM for Windows™ client GUI.

When you decide what files to include in a migration job, consider both the frequency of use of the files and the recall speed. Although most file recall is not noticed by users, network bandwidth, storage repository speed, and file size all determine the file recall speed.

A migration job file can be shared among computers with similar configurations, and can be shared among nodes in a cluster. If some directory structure of two computers is the same, you can use without modification a migration job that specifies the common directory structure on both computers.

Creating migration jobs

Use the HSM for Windows™ client GUI to define migration jobs. Migration jobs select different file sets to migrate by specifying different include and exclude conditions such as file age, size, subdirectory, and groups on files or directories.

About this task

The length of the path and name of migrated files is limited.

The length of a file name that is migrated by the HSM for Windows™ client cannot exceed 256 bytes. The path length (the API high-level qualifier) cannot exceed 1024 bytes. A path and file name includes the file server name, the volume, and the directory portion of the full Uniform Naming Convention (UNC) name, for example `\\FILESERVER\\E:\\directory\\filename.ext`. The Unicode representation of a character can occupy several bytes, so the maximum number of characters that a file name might contain can vary.

When using the HSM for Windows™ client GUI, path names can be a maximum of 254 characters only. For path names that exceed 254 characters, you must use the `dsmc1c.exe` command from a Command Prompt window.

To complete the following steps to define a migration job, run the HSM for Windows™ client GUI.

Procedure

1. Select **Job > New Job** or right-click in the window's white space and select **New Job**.

2. Name the new job icon to a name of your choice.
You cannot use a directory delimiter in the job name.
3. Double-click the new job icon to display the job creation window.
4. In the **General** tab, use the **File Space** menu to select the name of the file space in which you want to store migrated files.
5. In the **Backup before migration** box, you can specify that files must be backed up before they are migrated. If a job specifies a file that was not backed up, the file is backed up and then it is migrated. If you select this option, you must also indicate an options file for the backup before migration. You can specify an options file, or you can specify that the backup-archive client determines the options file.
6. In the **Management class** panel, select a management class for migrated files. A message at the bottom of the panel indicates the suitability of the management class for retaining the migrated files.
7. In the **Management class** tab, select a management class for migrated files. A message at the bottom of the panel indicates the suitability of the management class for retaining the migrated files.
8. On the **Migration Options** tab, specify whether file security attributes (ACL) and Windows™ alternate data streams (ADS) data are migrated when the file is migrated.
The defaults are the values that you set in the initial configuration wizard.
9. To add a directory, skip to step **"10" on page 44**. For each file you want to add, follow these substeps:
 - a. Select the Source Files tab's **New File** button.
 - b. Select **Browse**. In the **Browse for File** window, select the drive that you want and select **OK**.
 - c. Use the file selection window that displays to drill down to the file that you want and select **OK**.
 - d. Select a migration action. The default **Replace the file with a shortcut to the file space** option migrates the file and creates a stub file. The **Keep the original file** option migrates the file, but the original file remains on the local system. The **Delete the file** option migrates the file and then deletes it from the local system.

Note: Do not run reconciliation on the file spaces that are used for this job, if you select **Delete the file**.

- e. Select the **Source File** window's **Advanced Conditions** tab and select **New Include**. The following steps use the **Include Conditions** windows as examples, but you can also choose **New Exclude**, which follows the same convention. And you can combine include and exclude conditions.
 - f. From the **Include Condition** window's top menu, choose the type of condition you want for the selected files, define the settings, and select **OK**.
 - g. Continue to define include and exclude conditions for the selected files and select **OK** when complete.
10. To add directories from the **New Job** window's **Source Files** tab, select **New Directory** and then **Browse**. Select the directory that you want to add and select **OK**. Continue to add as many directories as required, then follow these substeps to define the details of the migration job:

Note: The migration action and include and exclude conditions that you apply to a subdirectory-based migration job applies to the individual files in the selected subdirectories.

- a. Select a migration action. The default **Replace the file with a shortcut to the file space** option migrates the file and creates a stub file. The **Keep the original file** migrates a copy of the file, but the original file remains on the file system. The **Delete the file** option migrates the file and then deletes it from the file system.

Note: Do not run reconciliation on the file spaces that are used for this job, if you select **Delete the file**.

- b. Select the **Include Subdirectories** check box if you want to include all files in the selected directory's subdirectories.
- c. Select the **Advanced Conditions** tab, and then select the type of include condition that you want to define.

Examples of including and excluding files

The following examples show the interaction of include and exclude conditions.

Note: The following examples are to help you get started with building your own include and exclude conditions. Test your own conditions thoroughly.

"Example base file set" on page 45 lists the base file set used in these include and exclude examples. A base file set includes all files in the selected disk, directories, and, if selected, all subdirectories. The content of the base file set never changes. Include and exclude conditions that you define create a subset of the base files that are valid for the selected operation. This valid subset of files is called the "target set". If you set no conditions, the HSM for Windows® client by default includes all files.

| Table 7: Example base file set | |
|--------------------------------|-----------|
| File name | File size |
| test.log | 1.5 GB |
| test.html | 50 K |
| test.bmp | 250 MB |
| test.pdf | 2.7 GB |
| test2.pdf | 11 GB |
| test.dwg | 100 GB |

Example 1: one include condition

```
include all files < 300 MB
```

The exclude condition is evaluated against all files in the base set. The result is all files that are less than 300 MB:

```
test.html (50 KB)
test.bmp (250 MB)
```

Example 2: one exclude condition

```
exclude all files < 300 MB
```

The exclude condition is evaluated against all files in the base set. The result is all files that are 300 MB or greater:

```
test.log (1.5 GB)
test.pdf (2.7 GB)
test2.pdf (11 GB)
test.dwg (8 GB)
```

Example 3: one exclude condition

```
exclude all files < 30 GB
```

The exclude condition is evaluated against all files in the base set. All files match the condition, so all files are excluded.

Example 4: two include conditions

```
include all files < 300 MB
include all files with extension = pdf
```

First, the first include condition is evaluated against all files in the base set. The result is the following files:

- test.html (50 KB)
- test.bmp (250 MB)

Next, the second include condition is evaluated against all files in the base set. The result is the following files:

- test.pdf (2.7 GB)
- test2.pdf (11 B)

The final result is any file that matches any of the include conditions:

- test.html (50 KB)
- test.bmp (250 MB)
- test.pdf (2.7 GB)
- test2.pdf (11 GB)

Example 5: two exclude conditions

```
exclude all files < 300 MB
exclude all files with extension = pdf
```

First, the first exclude condition is evaluated against all files in the base file set, and the result is the following files:

- test.log (1.5 GB)
- test.pdf (2.7 GB)
- test2.pdf (11 GB)
- test.dwg (8 GB)

Next, the second exclude condition is evaluated against the result of all previous evaluations. The final result is the following files:

- test.log (1.5 GB)
- test.dwg (8 GB)

Example 6a: mixed include and exclude conditions

This example coding does not yield a set of only PDF files that are less than 3 GB.

```
exclude all files < 3 GB
include all files with extension = pdf
```

First, the exclude condition is evaluated against all files in the base file set, and the result is the following files:

- test2.pdf (11 GB)
- test.dwg (8 GB)

Next, the include condition is evaluated against all files in the base file set. The result of the include condition is the following files:

- test.pdf (2.7 GB)
- test2.pdf (11 GB)

The final result is the following files:

- test.pdf (2.7 GB)
- test2.pdf (11 GB)
- test.dwg (8 GB)

Remember: An include condition is evaluated against all files in the base file set regardless of the preceding include or exclude conditions.

Example 6b: mixed include and exclude conditions

This example coding does yield a set of only PDF files that are less than 3 GB.

```
include all files with extension = pdf
exclude all files < 3 GB
```

First, the include condition is evaluated against all files in the base file set, and the result is the following files:

- test.pdf (2.7 GB)
- test2.pdf (11 GB)

Next, the exclude condition is evaluated against the set of files that result from all previous evaluations. The final result is the following files:

- test2.pdf (11 GB)

Example 7: redundant exclude condition

This example illustrates how an exclude condition can be redundant.

```
include all files with extension = html
exclude all files with extension = log
```

First, the include condition is evaluated against all files in the base file set, and the result is the following files:
test.html (50 KB)

Next, the exclude condition is evaluated against the set of files that result from all previous evaluations. The final result is the following files:
test.html (50 KB)

File groups

To facilitate the grouping of files for migration, you can create and edit file groups in HSM for Windows™. You define file groups by file extension types.

You can associate any number of file types to one file group. For example, you can have a group that is called "Image Files" consisting of these file extensions: bmp, jpg, eps, and gif. You can define another file group called "Office Files" consisting of the following file extensions: doc, xls, and ppt.

Note:

- A file group can be used in the definition of migration jobs.
- Every file group is global and any changes to the group changes its definition anywhere that group is used or selected.

Creating a new file group

Use these steps to create a new group using the HSM for Windows™ client GUI.

About this task

Note: The creation of a new file group is global. The new file type you create here will be included in the lists of types under **Tools > File Groups**.

Procedure

1. Select **Tools > File Groups**.
2. Click the **New file group** button.
3. Enter the name of the file group you want to define.
4. Enter the file extensions you want to be included in this file group, separated by spaces.
5. Click the **OK** button.

Edit a file group

Use these steps to edit an existing file group using the HSM for Windows™ client GUI.

About this task

Note: Any changes you make to a file group affect that file group globally, wherever it is used or selected.

Procedure

1. Select **Tools > File Group**.
2. Select the file group you want to edit and select the **Edit** button.
3. Edit the file extensions you want to be included in this file group.

Calculating a migration job's space savings

Before finalizing a migration job, you can calculate the amount of space that will be saved by a migration without having to run the migration job.

About this task

To calculate a migration job's space savings perform the following step:

Procedure

- Right-click the migration job that you want to calculate and select **Calculate Space Saving**.
The HSM for Windows™ client searches for all files that match the job criteria. If the file system contains many directories and files, the search can take some time. When all files have been searched, you can see three sets of information in both files count and kilobytes:
 - Current Disk Usage
 - Disk Usage after Migration
 - Free Disk space Gain

Migration jobs start by a schedule, GUI, or CLI

Migrations jobs can be started by a standard scheduler, by the HSM for Windows™ client GUI (graphical user interface), and by the HSM for Windows™ client CLI (command-line interface).

You can run migrations jobs any of the following ways:

- From the HSM for Windows™ client GUI
- From the **Command Prompt** window by using the **dsmc1c** command
- From a scheduled task

Running migration jobs from the HSM for Windows™ client GUI

After defining migration jobs, you can run them at any time from the HSM for Windows™ client GUI.

Run migration jobs from the HSM for Windows™ client GUI by right-clicking on a migration job and selecting **Execute Job Immediately**.

Viewing migration job results

When a migration job finishes, you can view the results.

About this task

When a migration job finishes, an information window displays.

Procedure

1. Click **OK**.
The **Task List** window opens.
2. Check the **Display per file details when migration is finished** box.
The detailed result is displayed when you close the **Task List** window.
3. Click **Report**.
The **Migration Report** window opens.
4. In the **Migration Report** window, click **Close**.
The **Migration Report** window closes.

5. In the **Task List** window, click the **Close** button.
The **Task List** window closes. The **Result** details window opens.
The **Result** window contains a list of the processed files and a message about the migration result for each file. Click the column headers to sort the **Name** and **Message** columns. Right-click a row to display information filters. Check or uncheck the filters to apply the filters to the list. The Show Stub Files filter is persistent and remains activated or deactivated until the status is changed by the user. The other three filters are activated by default and changes are valid only for the current GUI session.

Scheduling a migration job

You can schedule migration jobs to run automatically by using a scheduler provided by another vendor. Schedule the **dsmc1c.exe** command, specifying the job file as an argument when **dsmc1c.exe** is started.

About this task

You can run only one **dsmc1c.exe** process at a time. You cannot schedule two migrations at the same time, and you cannot schedule two migrations that overlap. The following steps show how to configure the Windows® Scheduler to start a migration job weekly:

Procedure

1. From **Windows Start** menu, select **Administrative Tools > Task Scheduler**.
The **Task Scheduler** window opens.
2. Click **Create Basic Task**.
The **Create Basic Task Wizard** window opens.
3. Type a task name and description. Click **Next**.
The **Trigger** window opens.
4. Click weekly (or as often as you want to run the task). Click **Next**.
The **Weekly** window opens.
5. Enter schedule details. Click **Next**.
The **Action** window opens.
6. Check **Start a program**. Click **Next**.
The **Start a program** window opens.
7. Type the path of the **dsmc1c.exe** command in the **Program/script** field. Type the job file name in the **Add arguments (optional)** field. Click **Next**.
The **Summary** window opens.
8. Click **Finish**.
Windows® creates the scheduled task.

Removing unused stubs from a file system

You can remove unused stub files from a file system by using a migration job. Migration copies are protected in IBM® Storage Protect until you decide that the migration copies are no longer required.

About this task

A file system can become populated with stubs of files that were migrated and were not recalled in a long time. Users can delete their obsolete files, but often they do not. An administrator can remove unused stub files from the file system and keep the migrated files in IBM® Storage Protect storage indefinitely.

Removing unused stub files from a file system has the following benefits:

- An administrator does not rely on a user to delete obsolete files.
- The administrator can choose which stub files are removed. The administrator can specify folders and age criteria to identify unused files.
- File system operations can be more efficient when fewer files exist on a file system. Stub files that are removed are not scanned during a file system scan.
- Listings of unused files do not distract from the listings of newer files. Stub files that are removed are not listed in the HSM for Windows™ client listing files.

- Unintended recalls and potential out-of-space conditions are minimized. If a user unintentionally copies a folder of obsolete files, the files must be recalled to the file system. If the unused files are not on the file system, the user cannot make such an error.
- A migration job that removes stub files does not retrieve or recall the migrated files to the file system.
- After the stub files are removed from the file system, the files are protected in IBM® Storage Protect storage. The files are not deleted from IBM® Storage Protect storage by running a standard reconciliation process. You can retrieve the files from IBM® Storage Protect storage by using the HSM for Windows™ client search and retrieve function.
- When migration copies are no longer required, the administrator can run a special reconciliation process that deletes the protected files from IBM® Storage Protect storage.

To remove unused stub files, complete the following steps.

Procedure

1. Determine the number of old stub files that are in a file system.
 - a. Run the **dshmsmc1c.exe** command with the **oldstub** parameter. Specify an age that defines an old stub file.
 - b. Run a reconciliation process for the file system.
The reconciliation process counts the number of stub files on the file system that are at least as old as the age that you specify.
 - c. View the **hsmmonitor.log** file.
The **hsmmonitor.log** file contains the number of stub files that are at least as old as the age that you specify. A trace record from the log file looks like this example:

```
I: Number of old/unused stubs (age > 400 days): 13467
```

You can repeat the process and use different ages. Use the information to determine the age of the stub files that you want to remove.

2. Create a migration job for removing unused stub files from the file system.
A job for removing unused stubs is similar to a job for migrating files, with the following caveats:
 - In the **General** tab, you must set the **Action** option to **Delete the files**.
 - When you create any migration job, you must select a file space and management class. However, migrated files are assigned to a file space and a management class when they are migrated. The file space and management class of the migrated file do not change when you remove a stub file from the file system. The file space and management class values are ignored when the job removes a stub file.
 - In the **Advanced Conditions** window, you must include a condition for migration status. Select **HSM stub file**.
 - In the **Advanced Conditions** window, you can exclude a condition for the stub file age or the time of the last migration.

Note: For files that were migrated with the HSM for Windows™ client 7.1.4 and earlier, the last migration time is set when the file is migrated and when the ACL of the stub file is updated. For files that were migrated with version 7.1.6 and later, the last migration time is set only when the file is migrated.

3. Run the HSM for Windows™ client **dsmfind** command. Specify the new migration job as a command parameter. Inspect the output list of files and decide whether the include and exclude conditions define the correct set of files to remove.
No stub files are removed from the file system when you run the **dsmfind** command.
4. Modify the new migration job include and exclude conditions and run the **dsmfind** command until the migration job defines the appropriate stub files to remove.
5. Run the migration job.
Stub files are removed from the local file system.

The migration copies of the files remain in IBM® Storage Protect storage. The migration copies are protected from standard reconciliation processes. The migration copies are deleted from IBM® Storage Protect storage only when you configure and run a reconciliation process to delete protected files.

What to do next

You can retrieve the files from IBM® Storage Protect storage by using the HSM for Windows™ client search and retrieve function.

Migration by file list

You can migrate a list of files that are contained in a text file. The text file can be created by any program, but must meet encoding and format criteria.

Migration jobs migrate files that meet a job's selection criteria. Threshold migration uses file size and age to determine which files to migrate, but you cannot specify which files are migrated. If you want to migrate specific files, regardless of age and size, you can do a list migration.

The list file must meet these specifications:

- The file is encoded in the Windows® default ANSI system code page or in Unicode. If the file is encoded in Unicode, it must be UCS-2LE, with a Byte Order Mark (BOM) as the first 2 bytes in the file. The BOM (0xFF, 0xFE) is automatically written when you save the file from a Notepad editor and specify Unicode encoding. UCS-2LE supports all languages that are supported by the HSM for Windows™ client.
- Each line of the file contains the complete path name of one file.
- Each line of the file is separated by carriage return and line feed (CRLF).

You can use another application to create the list file. Start the `dsmc1c.exe` command, specifying the `migratelistoption`, and specify the name of the list file.

Threshold migration

You can migrate files from your volumes according to high and low thresholds of space usage. With proper configuration, you can greatly reduce the chance of your volumes running out of space.

Threshold migration provides automatic control of space usage of the volume. You set the high and low space-usage thresholds that trigger the HSM for Windows™ client to automatically start and stop migration. You configure guidelines for migration candidates. HSM for Windows™ client uses those guidelines to choose which files to migrate, and when, to meet the space usage settings.

You can configure threshold migration with the **Threshold Migration settings** window in the GUI, or with the `dsmhsmc1c.exe` command.

Migration candidates

HSM for Windows™ client chooses larger and older files as candidates for threshold migration.

Files that are frequently modified or accessed are poor candidates for migration. HSM for Windows™ client assumes that the last access date or modification date or creation date is an indicator of how dynamic a file is. Hence, HSM for Windows™ client chooses migration candidates that have a greater age, as measured by access, modification, or creation date. You configure which of these dates (access, modification, or creation) HSM for Windows™ client uses to determine file age. You also configure the minimum age for a migration candidate. Among files that meet the minimum age, and are the same size, HSM for Windows™ client migrates only the oldest files.

Small files are also not good candidates for migration because migrating a small file frees up less space than migrating a large file. There is a transaction cost for every file migration and recall. The transaction cost is the same regardless of file size, even though migrating larger files frees up more space. Hence, HSM for Windows™ client chooses large files for migration candidates. You can configure the minimum size for a migration candidate, but among files with the same age, HSM for Windows™ client migrates only the largest files.

You can also configure the weight (importance) of age, relative to size, for migration candidates. For example, assume that your volume contains some large files that tend to be dynamic. You can decrease the chance that the files will be migrated by increasing the weight of file age for migration candidates.

To find the migration candidates, HSM for Windows™ client scans the volume. HSM for Windows™ client scans all directories in the volume in an orderly manner, but typically not all at once. A scan continues until enough migration candidates are found. The next scan starts where the previous scan finished, until, over time, the entire volume can be scanned. Further scans traverse the volume again and again. You can configure how often to scan for migration candidates.

If not enough migration candidates are found, HSM for Windows™ client can scan the entire volume in a single scan. If the entire volume is scanned without yielding sufficient candidates, HSM for Windows™ client issues a warning. At the next scan, it is possible that the size or age of some files will qualify them for migration.

The size and age of the files in the most recent scan are compared with the files in the migration pool. The comparison yields a new ranked list of migration candidates. The oldest and largest files are the first on the list.

A scan begins in these situations:

- The configured time interval since the last scan elapses.
- You manually start a scan.
- Before a threshold migration, the pool does not contain enough files to reduce the space usage from the high threshold to the low threshold.
- During a threshold migration, the pool of migration candidates becomes empty.

Migration candidates are stored in a pool, ready to be migrated when space usage reaches the high threshold. Before a migration, there must be enough migration candidates in the pool to reduce the space usage from the high threshold to the low threshold.

The pool contains more files than are needed, in case some candidates are no longer valid by the time of the next threshold migration. A file in the pool can lose migration eligibility for several reasons:

- The file was deleted from the file system.
- The file was modified, and it no longer meets the minimum age or the minimum size for migration.
- The configured minimum age or minimum size for migration was increased.

Periodically HSM for Windows™ client validates the files in the pool. Files that are no longer valid are eliminated from the pool. If the pool does not contain enough files to reduce the space usage from the high threshold to the low threshold, a scan begins. You can configure the frequency of the validation.

Migration triggers

Migration is automatically triggered when the HSM for Windows™ client detects that space usage reaches the high threshold. You can also start threshold migration manually, any time that space usage is greater than the low threshold.

The IBM Storage Protect™ HSM Monitor Service monitors space usage on an interval that you configure. Migration is triggered when the IBM Storage Protect™ HSM Monitor Service detects a high threshold of space usage, and continues until usage reaches the low threshold. The HSM for Windows™ client can decrease the interval when space usage approaches the high threshold. Nevertheless, if space usage increases rapidly and is not checked frequently enough, it is possible that space usage can exceed the high threshold before migration begins.

Configuring threshold migration

You can configure threshold migration with the graphical user interface (GUI). Files are automatically migrated from the volume when space usage reaches the configured threshold.

About this task

Complete the following steps to configure threshold migration by using the HSM for Windows™ client GUI.

Tip: You can also configure threshold migration by using the `dsmhsmc1c` command.

Note: Threshold migration requires free disk space to store the names of migration candidates. The space that is required depends on the number of migration candidates and length of the file names. If the files have long file names, about 10 MB of free disk space is required for every 5000 migration candidates. For short file names, less space is required.




Procedure

1. Start the HSM for Windows™ client GUI. Select **Tools > Threshold Migration**. The **Threshold Migration settings** window opens. If the volume is configured for threshold migration, the current configuration values are displayed in the fields.
2. Choose values for the threshold migration options and then click **OK**. The following threshold migration options and controls are available:

Mount path

Specify the volume mount path. Because it is possible for a single volume to be mounted by more than one path, always specify a volume by the same mount path. Reconciliation, threshold migration, and migration jobs must all reference the volume by the same path.

The icon indicates the status of the volume:

- Not configured:

- Configured:

- Not configurable:


The volume of this mount path is already configured through another mount path and cannot be configured through the path now selected.

Status

The field displays the current configuration status of the selected volume and whether a migration, scan, or validation process is running. Click **Refresh** to refresh the status.

Configure/Unconfigure button

When the volume is not configured, the button displays **Configure**. Click this button to activate the fields and controls in the window, and populate the fields with default values.

When the volume is configured, the button displays **Unconfigure**. Click this button to remove the configuration of the volume.

Management class

Use this option to configure the management class that is used for threshold migration of this volume. Specify an existing management class with an archive copy group, or specify **DEFAULT** to use the default management class of the active policy set. If the retention period of the selected management class is finite, a warning is issued.

Low threshold (%)

Use this option to configure the disk usage that triggers when to stop threshold migration. After the disk usage reaches this percent of capacity, threshold migration stops. The low threshold must be less than the high threshold. The range of acceptable values is 0-99. The default is 80.

High threshold (%)

Use this option to configure the disk usage that triggers when to start threshold migration. After the disk capacity reaches this percent of capacity, threshold migration begins. The range of acceptable values is 1-100. The default is 90.

Migrate to file space

Use this option to configure the file space that is used for threshold migration.

Back up files before migration

Use this option to configure whether a file must be backed up before it is migrated. The default is the value that you set in the initial configuration wizard.

Migrate file security (ACL) when a file is migrated

Use this option to configure whether file security attributes are migrated when the file is migrated. The default is the value that you set in the initial configuration wizard.

Migrate alternate data streams (ADS) when a file is migrated

Use this option to configure whether Windows™ alternate data streams data is migrated when the file is migrated. The default is the value that you set in the initial configuration wizard.

Select an IBM® Storage Protect options file for backup before migration

Use this option to specify the options file for backup before migration.

Space usage monitor interval (minutes)

Use this option to configure how frequently the HSM monitor service checks space usage on the disk. The time is measured in minutes. If the monitor interval is set to 0, monitoring is deactivated. The range of acceptable values is 0-9999. The default is 5.

Migration candidates scanning interval (hours)

Use this option to configure how frequently the HSM monitor service starts the file system scan to find migration candidates. The time is measured from the end of the last scan to the beginning of the next scan. The time is measured in hours. The range of acceptable values is 1-9999. The default is 24.

If a scan yields better quality candidates (older and larger files) than the previous scan, the interval is automatically decreased by a small amount. If a scan yields poorer quality candidates (newer and smaller files) than the previous scan, the interval is automatically increased by a small amount.

Migration candidates validation interval (minutes)

Use this option to configure how frequently the HSM monitor service validates the candidates in the candidates pool. The time is measured from the end of the last validation to the beginning of the next validation. The time is measured in minutes. If the interval is set to zero, validation is deactivated. The range of acceptable values is 0-9999. The default is 180.

Migrate now

Use this option to configure an immediate threshold migration. If disk usage is greater than the low threshold, files are migrated until the low threshold is reached. The default is no.

Scan now

Use this option to configure an immediate scan of the volume to find migration candidates. The default is no.

Minimum file size (KB)

Use this option to configure minimum file size for a valid migration candidate. The size is measured in kilobytes (KB). The range of acceptable values is 4-2147483647 (2 TB). The default is 4.

Minimum file age (days)

Use this option to configure minimum file age for a valid migration candidate. The age is measured in days. The range of acceptable values is 0-99999. The default is 360.

File age criteria

Use this option to configure which time stamp is used to calculate the age of a file. Changing this option can make many files in the current pool of migration candidates no longer valid. The choices correspond to the file system time stamps for file creation, file modification, and file access. The default is the file access time.

Weighting of age criteria (%)

Use this option to configure the importance of file age (relative to file size) when determining migration candidates.

The age weight and size weight of a file are computed relative to the configured minimum age and minimum size. Hence, a file that is twice as old as the minimum age has an age weight of 2. If the file is the minimum size, it has a size weight of 1.

When the importance of age relative to size is considered, the file's weight is computed in this way:
computed weight = (AGEWeight * (age weight)) + ((1-AGEWeight) * (size weight)).

For example, when AGEWeight = 50, the file has the same weight $((.5 * (2)) + ((1 - .5) * (1)) = 1.5)$ as a file that is only as old as the minimum age, but twice as big as the minimum size $((.5 * (1)) + (.5 * (2)) = 1.5)$. The weight of both files is 1.5.

If the AGEWeight option is not 50%, but 75%, the first file has a computed weight of 1.75 $((.75 * (2)) + ((1 - .75) * (1)) = 1.75)$, while for the younger but larger file, the computed weight is 1.25 $((.75 * (1)) + ((1 - .75) * (2)) = 1.25)$.

Specify a value from 0 to 100. The default is 50.

Maximum number of parallel threshold processes

Use this option to configure the number of migration tasks that can occur simultaneously. The option applies to migration, scan, and validation tasks on all volumes. If this number is reached, any pending migration tasks are delayed until one of the running tasks finishes. The range of acceptable values is 1-16. The default is 3.

Cleanup

When one or more configured volumes are no longer available, the **Cleanup** button is activated. Click this button to erase the configuration information for each of these volumes.

Refresh

Click **Refresh** to show the latest values. For example, if you added a file space since opening the window, click **Refresh** to show the current file spaces.

Apply

Click **Apply** to apply the configuration to the volume and leave the window open. Use **Apply** to reuse configuration setting when you configure several volumes.

OK

Click **OK** to apply the configuration to the volume and close the window.

Space management of the system volume

You can run migration jobs and list migrations on the Windows™ system volume. Do not configure threshold migration on the Windows™ system volume.

Attention: In threshold migration, files are migrated based on age and size. You cannot ensure that critical system files are not migrated. If you configure threshold migration on the system volume, it is possible that some critical files will be migrated. It is possible that the computer will become unusable or will not start.

If you run migration jobs or list migrations on the system volume, do not migrate critical system files.

Selectively retrieving and recalling migrated files

You can return selected migrated files to the originating file system. You do not have to wait for a file to be automatically recalled.

About this task

A file is recalled automatically when you or a Windows™ application accesses the stub file. You can manually return a migrated file to the file system by using the information on the IBM Storage Protect™ server or the information in stub files.

You retrieve migrated files by using the information in the IBM Storage Protect™ file spaces. If a migrated file exists in IBM Storage Protect™ storage, you can retrieve the file.

You recall migrated files by using the information in stub files on the file system. If a stub file exists on the file system, you can recall the file.

Restriction: You cannot use the IBM Storage Protect™ backup-archive client to retrieve files that were migrated by the HSM for Windows™ client.

Retrieving migrated files

Search the IBM® Storage Protect server file spaces to retrieve selected files.

About this task

If you configured the HSM for Windows™ client to keep or delete the original file on the file system, there is no stub. The migrated file is not automatically recalled when the resident file is accessed on the file system. You can access the migrated copies on the IBM Storage Protect™ server only by retrieving the files.

Tip: You can also use the HSM for Windows™ client `dsmc1c retrieve` command to retrieve migrated files.

Complete the following steps to search for and retrieve migrated files.

Procedure

1. Open the HSM for Windows™ client GUI.
2. Select **Migrate Retrieve > Search & Retrieve**.
3. Select values for the **IBM Storage Protect™ server** and **File Space** fields in which you want to search for files.
4. Specify your search criteria in the **Backend server query** tab and click **Search**.
If you renamed the stub file on the file system, the stub file name does not match the name of migrated file in the IBM Storage Protect™ file space. You must specify the name of the migrated file in the IBM Storage Protect™ file space.

If you do not specify at least one search criterion, all of the files that are stored in the file space are shown. The **Path** and **Filename** fields are case sensitive, but the **Volume** field is not case sensitive. You can use wildcards in any field. An asterisk (*) matches zero or more characters, and a question mark (?) matches a single character.

You can further refine your search results by using the **Result filters** tab. You can specify one or both of the following filters:

- **Migration action** is the action that is specified to the file found on the backend server. The migration action of files that are migrated with HSM versions earlier than version 7.1.4 is unknown.
- **Migration time** is the time when the backend file was last migrated. The migration time of files that are migrated with HSM versions earlier than version 7.1.6 is unknown.

The **Search Results** window opens.

5. Click **Select All** to retrieve all files or select individual files and then click **Retrieve**.
The **Retrieve options** window opens.
6. **Optional:** Choose a version to retrieve.
If you selected only one file, you can choose which version to retrieve. If you selected more than one file, the **Version** option is not available.
7. **Optional:** Indicate whether security information is retrieved.
If any of the selected files were migrated with security information, you can retrieve security information when the file is retrieved. If none of the selected files were migrated with security information, the security option is not available. If the security information is not retrieved, the retrieved file inherits the default security attributes of the file system to which it is retrieved.
8. **Optional:** Indicate whether Windows™ alternate data stream (ADS) data is retrieved.

If any of the selected files were migrated with ADS data, you can retrieve ADS data when the file is retrieved. If none of the selected files were migrated with ADS data, the option is not available.

9. Choose an option for overwriting files on the file system.
10. Click **Retrieve** to retrieve the selected files.

Selectively recalling migrated files

You can search a file system and selectively recall migrated files.

About this task

You can selectively recall only files that were replaced with stubs when they were migrated. You must recall files with an HSM for Windows™ client command. You cannot recall files with the HSM for Windows™ client GUI.

Restriction: When selectively recalling a file, you can recall only the primary data stream data. You cannot selectively recall the Windows™ alternate data stream (ADS) data.

Procedure

- From a DOS prompt, enter the HSM for Windows™ client **dsmc1c** command. Use the **dsmc1c recall** command to specify a single file path or a pattern with wildcards. Use the **dsmc1c recalllist** command to specify a file that contains a list of stub files. For example, to recall all migrated .xls files in the c:\projects\2013\ directory, issue the following command:

```
dsmc1c recall c:\projects\2013\*.xls
```

Automatic backup before migration

To protect your data completely, you must back up your data. The backup-before-migrate feature ensures that there exists a backup copy for every file that you migrate.

You can select whether the backup-before-migrate feature is used as the default option for all new migration jobs and threshold migrations. Use the **Backup before migration** window in the HSM for Windows™ client configuration wizard. In each job or threshold migration, you can accept that default or you can specify another option. For every job and threshold migration, you can choose whether to back up files before migration, and which options file to use for the backup. The backup-archive client automatically backs up the necessary files before migration. If the backup is successful, the file is migrated.

By default, the backup-archive client changes the access time stamp of a file when the backup-archive client backs up a file. If a migration job or threshold migration is configured to check a file's access time (-minagetype access), the file might not be migrated after a recent backup operation. Use the `preserveLastAccessDate` option of the backup-archive client to specify whether a backup operation changes the access time stamp.

Even if you schedule regular backups, a file can change and be migrated before a backup operation runs. At the next backup operation, the backup-archive client gets a copy of the migrated file. The migrated file is copied to the staging directory and is backed up. To avoid the file copies during a backup operation, use the backup-before-migrate feature and use the same backup options that are used for the scheduled backup.

The HSM for Windows™ client backup-before-migrate feature is not a substitute for regularly backing up your files. When you use the backup-before-migrate feature, the HSM for Windows™ client does not back up files in several cases:

- The HSM for Windows™ client does not back up stub files.
- The HSM for Windows™ client does not back up files that do not meet the migration criteria.
- The HSM for Windows™ client does not back up a file if there exists a current backup copy of the file.

Restriction: The path name length limits differ for migrated files and for files that are backed up before migration. When you back up files before migration, the file name is subject to the limitations of the backup-archive client. When you migrate a file, the file name is subject to the limitations of the API.

Choosing a backup options file

When files are backed up before migration, you can specify a backup options file, or you can let the backup-archive client determine the options file.

If you do not specify a backup options file for a backup before migration, the backup-archive client will determine the options file. The backup-archive client uses four methods to find an options file. The precedence of the methods is as follows:

1. An options file in a path specified by an environment variable
2. An options file in the directory from which the backup-archive client is invoked
3. An options file in the backup-archive client installation directory

If a file is regularly backed up with the backup-archive client default options file, then backing it up before migration with the backup-archive client default options file maintains a consistent set of backups. However, if a file is regularly backed up with an options file other than the default, you can specify this other options file for backups before migration. Using one options file for regular backups and a different options file for backups before migration can result in backup copies of the same file on two different IBM® Storage Protect servers.

If you specify a backup options file during the initial configuration of the HSM for Windows™ client, that options file is the default for all backups before migration. The backup-archive client does not determine the options file. You can specify different options files when you configure migration jobs and threshold migration. You can also specify a backup options file when you start migration using a HSM for Windows™ client command on the Command Prompt window.

Backup and restore of migrated files

Some types of backup can back up a stub or a complete migrated file. Six backup-archive client options control the backup and restore of migrated files.

The backup-archive client and the HSM for Windows™ client work together. The backup-archive client always maintains a copy of the resident file in the backup pool, whether this file is migrated or not.

In other words, for migrated files there are two identical versions of the file on the IBM® Storage Protect server. One version is in the HSM pool, created by the HSM for Windows™ client. And one version is the backup copy in the backup pool, created by the backup-archive client. When you restore files, the backup-archive client can always re-create the resident file from the backup copy, even if the copy in the HSM pool was deleted.

The **Skip migrated files** option and the **Check stub file reparse content** option regulate the backup of stub files. The two restore options, **Restore as migrated file**, and **Restore resident if not accessible**, define how migrated files are restored. The **Reset modified last access date** option determines whether the access time is changed when a file is backed up. The access time can affect migration. The **Staging Directory** option controls where copies of migrated files are temporarily stored by the backup-archive client.

There are some limitations for backing up migrated files:

- You must not use adaptive subfile backup and HSM. You must back up only the entire migrated file. If you use adaptive subfile backup on migrated files, you might not be able to restore migrated files correctly. The backup-archive client does not report any errors or warnings when you do a subfile backup of a migrated file.
- If `skipmigrated=yes`, the backup operation skips the migrated file. The stub is not backed up and the complete file is not backed up.
- If `skipmigrated=no` (the default), some backup types can back up the stub or the complete file. A backup has the following results:

Incremental backup or image backup

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

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

An incremental or image backup backs up the stub.

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

An incremental or image backup backs up the complete file.

Incremental-by-date backup

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

Selective backup or archive

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

Set the backup and restore options for migrated files in the backup-archive client options file **dsm.opt**. Use the preferences editor or directly edit the backup-archive **dsm.opt** options file. You can also specify an option when you start a backup-archive client command in a Command Prompt window.

Options for backing up of migrated files

Several options control how the IBM® Storage Protect backup-archive client backs up migrated files. A backup can skip migrated files, compare stub content, and use a temporary directory that you specify.

skipmigrated

When the **skipmigrated** option is set to **yes**, the backup-archive client does not back up or archive any stub files.

If the **skipmigrated** option is set to **no**, the backup-archive client can back up stub files during an incremental backup. The default value of the **skipmigrated** option is **no**.

checkreparsecontent

The value of **checkreparsecontent** is applied only when you specify the option **skipmigrated=no**. If you specify the option **checkreparsecontent=yes**, the backup-archive client compares reparse point content of the local stub file with the content in IBM® Storage Protect storage. If the content is the same, the stub file is not backed up again. If the local reparse point content is different from the backed-up content, the local stub file is backed up.

If you specify the option **checkreparsecontent=no**, the backup-archive client does not compare the reparse point content of the local stub file with the content in IBM® Storage Protect storage. Differences in the reparse point content are not detected, and no backup is created as a result of the reparse point comparison. If a valid stub file does not exist on the IBM® Storage Protect, you cannot restore a file as a stub file. In this case, you can restore a complete file instead of a stub file.

The **checkreparsecontent** option is one condition that can result in a file backup. Other conditions such as changes in file size or security settings are evaluated independently and can also result in a backup.

The reparse point of stub files that were backed up with the HSM for Windows™ client version 6.1 and earlier does not contain the same information as stub files that were backed up with later versions of the HSM for Windows™ client. As a result, all version 6.1 and earlier reparse points appear changed to later versions of the backup-archive client. If you set the option **checkreparsecontent=yes** and **skipmigrated=no**, the first incremental backup with a later-version backup-archive client creates new backup copies of all version 6.1 and earlier stub files. The new backups in IBM® Storage Protect storage contain the later-version reparse point information. Subsequent incremental backups create new backup copies of stub files only if the reparse point indicates that the file was changed.

When you set this option, IBM® Storage Protect checks the reparse point content of the local stub file, which increases the time for a backup operation. Set this option the first time that you do an incremental backup after either of the following events:

- You move migrated files with the **dsmmove.exe** command.
- You change the file space that is used for migration.

Clear this option on subsequent backups.

| Table 8: Interaction of options skipmigrated and checkreparsecontent during incremental backup | | |
|--|-------------------------------|--|
| | skipmigrated=yes | skipmigrated=no |
| checkreparsecontent=no | A stub file is not backed up. | A stub file is not backed up if only the reparse-point content changed. A stub file can be backed up if other changes occurred. |
| checkreparsecontent=yes | A stub file is not backed up. | Reparse point content of the local stub file is compared with content in IBM® Storage Protect storage. A local stub file is backed up if the content does not match. Also, a stub file can be backed up if other changes occurred. |

stagingdirectory

The backup-archive client ensures that whenever a stub is backed up, there is a copy of the complete file in the backup pool. If a complete file was not backed up before migration, the migrated copy is temporarily copied back and is backed up. IBM® Storage Protect associates the backup copy of the complete file with the backup copy of the stub. After the complete file is backed up, the temporary file is removed by the backup-archive client.

You can control the location to which the backup-archive client copies the temporary file by using the **stagingdirectory** option of the backup-archive client. When you use a staging directory for the temporary copy, the stub is not changed. The next backup creates a backup copy of the stub file on the IBM® Storage Protect server in the backup pool.

If the backup-archive client cannot create a complete backup copy of the migrated file, the backup-archive client does not back up the stub file. For example, if the stub is an orphan with no migrated copy in IBM® Storage Protect storage, the stub is not backed up.

IBM® Storage Protect maintains a backup copy of both the complete file and the stub. The backup copy of the complete file does not expire until the backup copy of the stub expires. Either the complete file or the stub can be re-created by using the backup-archive client.

If you set **skipmigratedno**, files that were not backed up before migration are copied to the staging directory when they are backed up. Many files are copied during a backup in the following situations:

- You have many stubs that were backed up with backup-archive client version 5.4 and earlier versions. The files are temporarily copied to the staging directory during backups with backup-archive client version 6.1 and later.
- You changed the backup policies for a volume by including for backup many files that were not previously included.
- You renamed stubs or directories that contain stubs.
- You changed the security settings of stubs or directories that contain stubs, and you configured migration of security attributes.

Managing backup performance when stub file encryption changes

You can limit the impact to backup performance that is caused by changing the encryption of stub files. Stage the encryption changes and backup operations.

About this task

If you change the encryption of a stub file, the IBM® Storage Protect backup-archive client copies the migrated file to a staging directory during the next incremental backup operation. If you change the encryption of many files, the backup operation can take a long time for the many temporary file copies.

To avoid temporarily copying many files, set the encryption status of files before you back up the resident files.

If you must change the encryption of an HSM-managed volume, you can stage the encryption change and the backup operation.

Tip: Another option is to first change the encryption of all the files. Then, back up migrated files separately from resident files.

Procedure

1. Change the encryption of the files in one directory of a volume.
2. Run an incremental backup of the changed files in the directory. Stub files with a modified encryption status are temporarily copied and backed up.
3. Repeat steps 1 and 2 for each directory in the volume.

Backing up migrated files separately from resident files

In some cases, you can limit the impact to backup performance by backing up migrated files separately from resident files.

About this task

The backup-archive client must temporarily copy a migrated file during an incremental backup operation if there is no backup copy, or if the stub file encryption changed. The temporary copies of migrated files can affect the performance of a backup operation. You can limit the impact to performance by backing up migrated files separately from resident files. Use the `skipmigrated` option to exclude migrated files from a backup operation.

When the `skipmigrated` option is set to `yes`, the backup-archive client skips migrated files. Permanently skipping the backup of migrated files can prevent you from recovering your data in a disaster. Temporarily skipping migrated files can reduce the time that is required for backup operations.

You can complete backup operations in a reasonable time and protect all your files if you run two types of backup operations. Run one backup operation only on resident files and run one backup operation on all files (resident and migrated files). The two backup operations protect all files.

Procedure

1. For your regular backup operations, set `skipmigrated=yes`.
Migrated files are excluded from regular backup operations.
The following backup-archive client command runs an incremental backup that skips migrated files:

```
dsmc inc N:\budgets\ -skipmigrated=yes
```

2. Run another backup operation with `skipmigrated=no`.
Files that are excluded from your regular backup operation are included. The backup operation makes temporary copies of the migrated files and can take a long time.
3. When the backup in [“2” on page 61](#) is complete, set `skipmigrated=no` for your regular backup operations.
The quantity of migrated files that must be temporarily copied is reduced from [“2” on page 61](#). All files (resident and migrated) are backed up in your regular backup operations.

Result

All files are backed up. Each backup operation completes in a reasonable time.

Options for restoring migrated files

Use the **Restore as migrated file** (`restoremigstate`) and **Restore resident if not accessible** (`restorecheckstubaccess`) backup-archive client options to manage how the backup-archive client restores migrated files from IBM® Storage Protect storage.

For files that are backed up with the backup-archive client, there is a backup copy of a resident file for every corresponding stub file. With the backup-archive client, you can restore the stub file or the resident file.

There are times when the IBM® Storage Protect HSM pool does not contain a copy of the migrated file, as shown in the following scenario:

1. A resident file is migrated to the IBM® Storage Protect HSM pool. A stub file remains on the volume.
2. The stub file is backed up. There is a backup copy of the stub file and a backup copy of the resident file in the IBM® Storage Protect backup pool.
3. The stub file is deleted from the volume.
4. During reconciliation, the migration copy in the IBM® Storage Protect HSM pool is deleted.

In this case, restoring the stub file can lead to problems because the HSM for Windows™ client cannot recall the migration copy of the file. If there is no migration copy in IBM® Storage Protect HSM pool, it would be better to restore the resident file rather than restore the stub. The backup-archive client can check whether a migration copy exists before restoring a stub file. If a migration copy does not exist, the backup-archive client can automatically restore the resident file instead of the stub file.

The **Restore resident if not accessible** (restorecheckstubaccess) and **Restore as migrated file** (restoremigstate) options configure how migrated files are restored by the backup-archive client. The options yield the restore results that are described in “[Results of using restoremigstate and restorecheckstubaccess options.](#)” on page 62

This table shows the results of using restoremigstate and restorecheckstubaccess options.

| <i>Table 9: Results of using restoremigstate and restorecheckstubaccess options.</i> | | |
|--|--|--|
| restorecheckstubaccess value | restoremigstate=no | restoremigstate=yes (the default) |
| restorecheckstubaccess=no | Restore the resident file; do not restore the stub | Restore the stub. Do not check whether a migration copy exists. |
| restorecheckstubaccess=yes (the default) | Restore the resident file; do not restore the stub | If a migration copy exists in the HSM pool, restore the stub. If a migration copy does not exist in the HSM pool, restore the resident file from the backup copy pool. |

In addition to the preceding options settings, the following conditions must also be true to restore a stub:

- The file was migrated at the time of the last backup
- The HSM for Windows™ client is installed
- The stub backup copy is an active version backup.
- The original file system and the target file system are of the same type (NTFS or ReFS)
- The stub is restored to the same path, and the file space name matches the volume name

There are some advantages to restoring a stub without checking that a migration copy exists in the HSM pool:

- Less temporary space is needed during restore
- There is less network traffic during a restore
- The restore is faster

There is a disadvantage to restoring a stub without checking that a migration copy exists in the HSM pool. There might be no migration copy in the HSM pool. If you restore a stub for which there is no migration copy, you create a stub file orphan. However, you can use reconciliation to report the stubs that are orphans. Then, you can restore the resident files from the backup pool with the option restoremigstate=no. If you run reconciliation in emulation mode, the HSM for Windows™ client creates a list of orphan stubs, but does not delete any files from IBM Storage Protect™ storage.

In the following examples, N:\file.txt was migrated, and a stub file remained on the volume. The stub file was backed up with the backup-archive client. Both the stub file and the resident file are available to the backup-archive client. The migrated file is restored by the backup-archive client with the **restore** command.

Task

Restore the resident file N:\file.txt.

Command: dsmc rest N:\file.txt -restoremigstate=no

Task

Restore a stub file N:\file.txt, regardless of whether a migration copy exists in IBM® Storage Protect HSM pool.

Command: dsmc rest N:\file.txt -restoremigstate=yes -restorecheckstubaccess=no

Task

Restore a stub file N:\file.txt, if a migration copy exists in IBM® Storage Protect HSM pool. If a migration copy does not exist in IBM® Storage Protect HSM pool, restore the resident file.

Command: dsmc rest N:\file.txt

Because the default option values are -restoremigstate=yes and -restorecheckstubaccess=yes, it is not necessary to specify the options.

Restriction:

- If the HSM for Windows™ client is not installed, or if the IBM Storage Protect™ HSM Recall Service is not running, default security attributes are applied to restored files.
- If a backup-archive client restore process is stopped in an unusual way (for example by pressing Ctrl+C or by restarting your system), files might remain in a temporary subdirectory (\~tsmtemp\) in the volume root. In this case, you must manually delete the \~tsmtemp\ directory.

Reconciliation

Reconciliation synchronizes your file system with the IBM® Storage Protect server by logging orphan stubs and by deleting obsolete copies of files.

You can use the HSM for Windows™ client graphical user interface (GUI) and the **dsmhsmc1c.exe** command to both configure and start reconciliation. You can start reconciliation at any time and can define reconciliation to run automatically in defined intervals.

The two main advantages of reconciliation are to reduce costs and to maintain integrity of your file systems. Reconciliation can reduce your costs by removing unnecessary or obsolete migrated objects from the IBM Storage Protect™ server storage. With fewer files, you require less storage. You also require fewer licenses because the HSM for Windows™ client is volume-licensed based on the amount of storage space that is used for migrated data on the IBM Storage Protect™ server.

The HSM for Windows™ client helps you maintain the integrity of your file systems by finding orphan stubs. Orphan stubs are stubs for which there is no migrated copy in IBM Storage Protect™ storage. Those orphans are recorded in the `hsmmonitor-orphan.log`. When you check the log file, you decide whether you want to delete the orphan stub or restore the stub from a backup.

If the reconciliation process finds any orphan stubs, the reconciliation process does not delete any object from IBM Storage Protect™ storage until all orphans are resolved. Resolve orphan stubs either by deleting the stub from the volume or restoring the complete file backup version.

If you delete a file but do not empty the **Recycle Bin**, the reconciliation process finds the file in the recycle bin. The reconciliation process does not delete the migrated copy from IBM Storage Protect™ storage.

Reconciliation deletes objects on the IBM Storage Protect™ server. For maximum data protection, back up all migrated files before you start reconciliation.

An object is deleted only after two reconciliation processes are run. After a stub is deleted, the first reconciliation marks the object on the IBM® Storage Protect server. If the stub is not restored before the second reconciliation, the object is deleted from IBM® Storage Protect server storage. If the stub is restored after the first reconciliation, the object is unmarked on the IBM® Storage Protect server and is no longer a candidate for deletion.

In previous versions, the reconciliation log file listed files that are deleted during reconciliation. In the current version, the reconciliation log file also lists objects that are marked or unmarked during reconciliation.

Reconciliation supports files that are migrated and replaced with stubs on the file system. Reconciliation is not intended for file spaces or volumes with migration jobs that have the action **Keep the original file** or **Delete the file**.

If files were migrated before a file system image backup was created, the file system image backup can contain stub files. After the image backup, the files can be recalled, and reconciliation can expire migration copies on the IBM Storage Protect™ server. When you restore the file system image, there can be stub files for which there

are no corresponding migration copies on the IBM Storage Protect™ server. In this case, you can restore an orphaned stub with the backup copy of the file that was created before the file was migrated.

If files were migrated after a file system image backup was created, the IBM Storage Protect™ server can contain migration copies for which there are no stub files. You can restore the stub files after the file system image restore. Restore the stubs before you run reconciliation. If you run reconciliation before you restore the stub files, the migration copies are deleted from the IBM Storage Protect™ server. Restoring stubs after the migration copies are deleted from the IBM Storage Protect™ server leaves orphan stubs on the file system.

A reconciliation process logs the actions that are taken against the objects on the IBM Storage Protect™ server. The log file is in the directory that is specified in the HSM GUI in the **Tracing Preferences** menu. The file name is `hsmmonitor-delete-YYYYMMDD-hhmmss.log`, where `YYYYMMDD` indicates the date and `hhmmss` indicates the time when the HSM monitor service was started.

If you run the reconciliation process in emulation mode (**reconcilemode=emulation**), the log file shows what actions would be taken in normal mode.

If the reconciliation process is run in normal mode (**reconcilemode=normal**), the list file contains the name of the obsolete objects. The objects are deleted by the reconciliation process. Normal mode is the default.

Reconciliation uses the name of the volume and the name of the nested volumes to identify files that do not belong to the file system.

If you rename a volume after migrating files, you must create a hardware volume mapping. If you do not create a hardware volume mapping, the reconciliation process can erroneously assume that the files were deleted from the file server. The reconciliation process can delete the files from the IBM Storage Protect™ server. If this situation occurs, use the backup-archive client to restore the complete file space to the renamed volume.

Tip: To improve reconciliation performance and avoid restoring files with the backup-archive client, use separate file spaces for each file system.

Changed volume mount-paths

If you change a volume mount point or drive letter or change the file server name, HSM for Windows™ reconciliation can be affected. You can mitigate many problems by creating hardware volume mappings, and prevent some problems by using unique file space names.

If you change a volume drive letter, mount point, or file server name, you can maintain HSM services by creating a hardware volume mapping. If you do not create a hardware volume mapping, reconciliation processing can delete migration copies in IBM® Storage Protect storage. If the drive letter, mount point, or file server name does not match the information on the IBM® Storage Protect server, a reconciliation process marks a migration copy as obsolete. The obsolete object is deleted from IBM® Storage Protect storage, subject to retention policy. A hardware volume mapping matches the old drive letter, mount point, or file server name with the new drive letter, mount point, or file server name. With a hardware volume mapping, a reconciliation process does not delete migration copies only because the drive letter, mount point, or file server name is changed.

A hardware mapping maintains HSM services for some changes to nested volumes. If you change only the drive letter or file server name, a hardware mapping continues HSM services. If you move a nested volume to another volume, a hardware mapping does not continue HSM services.

For example, assume that volume `\\MYNODE\E$\nested` is mounted into volume `\\MYNODE\E$`. Files are migrated from both volumes and reconciliation is configured for both volumes.

Assume that you change the drive E to F. Volume `\\MYNODE\E$` is renamed to `\\MYNODE\F$` and volume `\\MYNODE\E$\nested` is renamed to `\\MYNODE\F$\nested`. In this case, a hardware volume mapping continues HSM services, including accurate reconciliation.

Assume that you do not change drive letter E, but you move the nested volume into `\\MYNODE\G$`. The nested volume becomes `\\MYNODE\G$\nested`. In this case, a hardware volume mapping cannot maintain HSM services.

The second case can be mitigated with some planning. You can migrate the files of each volume to a separate file space on the IBM® Storage Protect server. Reconciliation can then be limited to only this file space. In this case, the IBM® Storage Protect server query, which is performed at the beginning of the reconciliation for a volume, does not return any objects from other volumes. The IBM® Storage Protect server does not delete any objects in storage that are from other volumes.

Tip: You can manage which file spaces are used during reconciliation with the **FILESPacelist** option of the **dsmhsmc1c** command or by using the **Reconcile settings** window of the HSM for Windows™ client GUI.

If a reconciliation process deletes objects from IBM® Storage Protect storage, you can restore the files from backup copies that were created by the backup-archive client. You can restore the complete file, even if the migration copy was deleted from IBM® Storage Protect storage.

Configuring reconciliation with the graphical user interface

Configure reconciliation with the graphical user interface (GUI) by using the **Reconcile settings** window.

About this task

Access the **Reconcile settings** window by selecting HSM for Windows™ client GUI. Select **Tools > Reconciliation**.

The **Reconcile settings** window displays configuration information. If the volume is not configured, the fields display default values. If the volume is configured, the fields display the current configuration.

Mount path

Specify the volume mount path. Because it is possible for a single volume to be mounted by more than one path, always specify a volume by the same mount path. Reconciliation, threshold migration, and migration jobs must all reference the volume by the same path.

Status

The field displays the current configuration status of the selected volume and whether a reconciliation process is running. Click **Refresh** to refresh the status.

Configure/Unconfigure button

When the volume is not configured, the button displays **Configure**. Click this button to activate the fields and controls in the window, and populate the fields with default values.

When the volume is configured, the button displays **Unconfigure**. Click this button to remove the configuration of the volume.

Next reconcile

Use this option to change the time of the next reconciliation. The field displays the date and time of the next reconciliation. If reconciliation is not configured, the default is the current date and time. If reconciliation is configured, the field displays the date that is calculated by adding the **Reconcile interval (hours)** to the last reconciliation.

Reconcile interval (hours)

Use this option to configure the number of hours between reconciliations. The interval starts when a reconciliation ends. If this option is set to 0, automatic reconciliation is deactivated. The range of acceptable values is 0-876000. The default is 720 hours.

Reconcile now

Use this option to reconcile the volume immediately. This action does not affect the **Reconcile interval (hours)** or the **Next reconcile** date.

File spaces used to reconcile

Use this option to configure the file spaces that are used during reconciliation.

You can improve the reconciliation performance by restricting the list to the file spaces that contain migrated files of the volume that you are configuring.

Remote IBM® Storage Protect Server Connections used for reconcile

Specify which remote IBM® Storage Protect server connections are used for reconciliation. By default, no remote IBM® Storage Protect server is included in reconciliation. If you select a remote IBM® Storage Protect server, all file spaces of the remote IBM® Storage Protect server connection are included in the reconciliation process.

If a file is recalled when in moving state, the migrated object is not automatically deleted on the remote IBM® Storage Protect server. The migrated object remains on the remote IBM® Storage Protect server until the remote IBM® Storage Protect server is added to a reconciliation process, and the reconciliation process is run.

Reconcile protected files

Set this option to reconcile protected files. A protected file is a file that was migrated and the file or stub file was deleted from the file system by a migration job. The default is to not reconcile protected files.

When you set the **Reconcile protected files** option, you can specify a time period. Specify the time period as a number of days. The reconciliation process processes protected files that became protected only before that time period. The default is 1095 days.

Maximum number of parallel reconcile processes

Use this option to configure the number of reconciliation tasks that can run at the same time. If this number is reached, any additional reconciliation tasks are delayed until the running reconciliation task finishes. Specify a value from 1 to 16. The default is 3.

Cleanup

When one or more configured volumes are no longer available, the **Cleanup** button is activated. Click this button to erase the configuration information for each of these volumes.

Refresh

Click **Refresh** to show the latest values. For example, if you added a file space since opening the window, click **Refresh** to show the current file spaces.

Apply

Click **Apply** to apply the configuration to the volume and leave the window open. Use **Apply** to reuse configuration setting when you configure several volumes.

OK

Click **OK** to apply the configuration to the volume and close the window.

Space requirements for reconciliation

Reconciliation uses Windows™ Volume Shadow Copy Service (VSS) to scan a volume. In addition to the VSS snapshot, VSS requires free disk space for the reconciled volume.

VSS requires space on the volume that is reconciled even if a snapshot is stored on another volume. VSS requires 200 KB of free disk space as a base requirement. Additionally, approximately 10 MB of disk space is required for every 100,000 objects on the reconciled file system.

VSS requires space for a snapshot. The snapshot can be on the reconciled volume or on another volume. Use the **vssadmin add shadowstorage** command to specify the volume for the snapshot. For information about the **vssadmin add shadowstorage** command, see *Vssadmin add shadowstorage* at the Microsoft™ technical notes library: technet.microsoft.com.

Previewing files that would be deleted by a reconciliation process

You can create a list of files that would be deleted by a reconciliation process. When you run a reconciliation process in the emulation mode, the files are not deleted.

Use the **reconcilemode** option with the **dsmhsmc1c** command to create a list file of obsolete objects on the IBM® Storage Protect server. When you specify the option **reconcilemode=emulation**, the reconciliation process does not delete obsolete objects, but writes the file names to the list file `hsmmonitor-delete-YYYYMMDD-hhmmss.log`. YYYYMMDD indicates the date and hhmmss indicates the time when the HSM monitor service was started.

Deleting protected files from IBM® Storage Protect storage

You can configure a reconciliation process to delete protected files from IBM® Storage Protect storage.

About this task

A protected file is a file that was migrated to IBM® Storage Protect storage and the file was deleted from the file system. Some migration jobs delete the file from the file system on the initial migration. Some migration jobs delete the stub file of a migrated file. Both kinds of jobs yield files that are protected in IBM® Storage Protect storage. A reconciliation process with default configuration values does not delete protected files. To delete protected files from IBM® Storage Protect storage, you must configure a reconciliation process with the option to delete protected files.

Procedure

1. Configure a reconciliation process to delete protected files.
A reconciliation process for deleting protected files is similar to a reconciliation process for unprotected files, with the following caveats:
 - In the **Reconcile Settings** window, you must set the **Reconcile protected files** option.
 - When you set the **Reconcile protected files** option, you can specify a time period. Specify the time period as a number of days. The reconciliation process processes protected files that became protected only before that time period. The default value is 1095 days.
2. You can test the configuration by running the **dsmhsmc1c** command with the `reconcilemode=emulation` option.
3. Run the reconciliation process.
Files that became protected before the time period are marked for deletion.
4. Run the reconciliation process again.
On the second reconciliation, the files are deleted from IBM® Storage Protect storage.

Moving migrated files

You can move migrated files to another volume on the same computer or to a volume on another file server.

You can move migrated files to accommodate the changing needs of users, applications, and hardware. For example, if a user moves to another site, you can move the migrated data. If a new or changed application requires that data is moved to another location, you can move the migrated files. You can maintain HSM services without recalling and migrating the files again.

If you do not plan the movement of migrated files, you can encounter several problems:

- Stub files can become inaccessible
- Many migrated files can be recalled, resulting in out-of-space conditions
- Tapes can be mounted several times

You can move stub files to another location with the **dsmmove** command. If the new location is managed by a different IBM® Storage Protect server, the HSM for Windows™ client moves migrated file data from the old to the new IBM® Storage Protect server.

The computer from which stub files are moved away is called the *remote* file server. The stub files on the remote file server are remote stub files. The IBM® Storage Protect server that manages the remote stub files is the remote IBM® Storage Protect server.

The computer to which the stub files are moved is called the *local* file server. The stub files on the local file server are local stub files. The IBM® Storage Protect server that manages the local stub files is the local IBM® Storage Protect server.

Migrated data is automatically moved when stub files are moved

If you move stub files to a location that is managed by a different IBM® Storage Protect server, the HSM for Windows™ client automatically moves the migrated data to the new server.

The **dsmmove** command uses the **hsmtasks** service on the local computer to complete the following tasks:

- Move the migrated data from the remote IBM® Storage Protect server to the local IBM® Storage Protect server. Only the version of migrated data that corresponds to the moved stub file is moved. The data is copied directly from one IBM® Storage Protect server to the other. No data is recalled to the file system.

- Remove the migrated data from the remote IBM® Storage Protect server, subject to the constraints of the retention policy. Only the version of migrated data that corresponds to the moved stub file is removed. Other versions of migrated data remain on the remote IBM® Storage Protect server. Other versions might belong to other stub files on the remote file server.
- Change the reparse content of the local stub file to point to the local IBM® Storage Protect server.
- Remove the *moving* state flag from the reparse content of the local stub file.
- Write a list file to the *installation_path*\listings directory. The file documents the movement of the migrated data between IBM® Storage Protect servers.

If the remote file system and the local file system are managed by the same IBM® Storage Protect server, the migrated data does not move to another IBM® Storage Protect server.

Stub files in moving state

Stub files can be in moving state until the move is complete. There are restrictions for stub files in the moving state.

If the remote file system and the local file system are not managed by the same IBM® Storage Protect server, the HSM for Windows™ client moves the migrated data to the local IBM® Storage Protect server.

Until the migrated file data is moved to the local IBM® Storage Protect server, the stubs are in *moving* state. The moving state is indicated by a flag in the reparse content of the local stub file.

The HSM for Windows™ client GUI indicates the status of the hsmtasks service.

When a migrated file is in moving state, you can search, retrieve, or delete it on the remote IBM® Storage Protect server. You can include a remote IBM® Storage Protect server for a reconciliation process.

If a remote stub file in moving state is recalled, retrieved, or renamed before it is moved, the stub cannot be moved. The HSM for Windows client creates a list of stubs that were not moved in the \tasks\error\ directory. Before you delete a connection to a remote server with the HSM for Windows client GUI, a warning message reminds you of the list.

A stub cannot be moved again while it is in the moving state, even if the move is on the same file server.

A stub in moving state depends on the current settings of the corresponding connection, which is stored in the `dsm.opt` options file in the HSM client installation directory. If you change any of the options in the `dsm.opt` options file, the stub in the moving state can no longer be accessed.

Moving stub files to another location

You can move stub files to another location. If the other location is managed by another HSM for Windows™ client or another IBM® Storage Protect server, the migrated data on the IBM® Storage Protect server is also moved.

Before you begin

You can move stub files to another location on the same file server and volume. You can move stub files to another volume on the same file server or on a different file server.

The location to which you move the stub files must be managed by an IBM® Storage Protect server. The IBM® Storage Protect server must provide HSM services for the location. You can move the stub files only to an NTFS file system or an ReFS file system.

The location from which you move the stub files must be managed by an IBM® Storage Protect server. This IBM® Storage Protect server is required until the move is complete.

All HSM for Windows® clients that are involved in the move must be IBM® Storage Protect Version 6.3 or later.

About this task

To move the stub files complete the following steps:

Procedure

1. Define the connection parameters for the remote IBM® Storage Protect server. Use the HSM for Windows™ client that is running on the local file server.

If stub files are moved to a clustered system, you must configure a connection on each node of the cluster. Configuring a connection for each node ensures that stubs that are in *moving* state can be accessed after failover.

- a. In the HSM for Windows™ client GUI, click **Menu > Tools > Remote IBM® Storage Protect Servers**.
- b. Click **Create**.
The remote connections wizard opens.
- c. From the remote connections wizard, enter the connection information in the wizard panels.
If stubs are moved from a clustered system, use the cluster name.

You must grant proxy authority to the remote node and use the **asnodename** option on the remote HSM nodes.

The remote connection wizard tests the connection. If the connection is successful, the HSM for Windows™ client creates a new options file in the `\config\` directory of the HSM for Windows™ client installation directory. The file name is constructed from the unique connection pair of server and node, and is file type `.opt`. An example configuration file name is `\config\server1-node1.opt`.

2. Move the stub files by using the **dsmove** command. Run the **dsmove** command on the local file server. The **dsmove** command moves the stub files to the local file system. If the local file system is managed by a different IBM® Storage Protect server, the **dsmove** command moves migrated data to the new server.

Continuing HSM services when a volume or file server is renamed

You can replace or rename the file server host and storage volumes. To continue HSM services, map the new volumes to the old volumes.

The HSM for Windows™ client uses the file server host name and drive letters. This information is used to identify the migrated object on the IBM® Storage Protect server during recall processing. If you change the drive letter of a volume that contains migrated files, IBM® Storage Protect cannot retrieve files. If you change the file-server host name or cluster name, IBM® Storage Protect cannot recall or retrieve files. If you change the drive letter, host name, or cluster name, a reconciliation process might mark the migrated objects on the IBM® Storage Protect as obsolete.

You must map the new volume to the old volume in the following situations:

- You rename the volume drive letter or mount point on a file server
- You replace the file server hardware or change the file server host name or cluster name

Renamed volume drive letters or mount points

You can map a volume drive letter. Any Universal Naming Convention (UNC) path within the mapped drive letter is automatically mapped. Volumes that are nested within the mapped drive letter are automatically mapped.

You cannot create a mapping for an individual nested volume. If you change the mount point of a nested volume, you cannot create a mapping for this individual mount point. You must create a mapping for the underlying drive letter. A mapping of the new drive letter to the old drive letter continues HSM services for some moves of nested volumes, but not for all moves. If you move a nested volume to another volume, a hardware mapping does not continue HSM services.

Example

For example, assume that volume `\\MYNODE\E$\nested` is mounted into volume `\\MYNODE\E$`. Files are migrated from both volumes and reconciliation is configured for both volumes.

Assume that you change the drive E to F. Volume `\\MYNODE\E$` is renamed to `\\MYNODE\F$` and volume `\\MYNODE\E$\nested` is renamed to `\\MYNODE\F$\nested`. In this case, a hardware volume mapping continues HSM services, including accurate reconciliation.

Assume that you do not change drive letter E, but you move the nested volume into `\\MYNODE\G$`. The nested volume becomes `\\MYNODE\G$\nested`. In this case, a hardware volume mapping cannot maintain HSM services.

New file server hardware or a changed file server host name

If you replace or rename a file server, you can attach space-managed volumes from the original file server. To continue HSM services, you must map the volumes on the new system with the volumes on the original system.

Example

Assume that an old file server is replaced by a new file server. The disk drives from the old file server are connected to the new file server. The new file server can have a different name, IP address, IBM® Storage Protect node name, and drive letters for the disk drives. If you map the volumes on the new system with the volumes on the original system, you can continue HSM services.

Hardware volume mappings

Hardware volume mappings are stored on the IBM® Storage Protect server in a private file space. The private file space requires a management class that does not expire objects. Changes to the IBM® Storage Protect server can affect the hardware volume mappings in the following ways:

- When an IBM® Storage Protect server database is restored, the mappings revert to the level of the restored database.
- If the IBM® Storage Protect server is changed, you must export and import the data in the private file space.

If a management class that does not expire objects is not available, the hardware volume mapping cannot be saved on the IBM® Storage Protect server. Not saving the mapping on the IBM® Storage Protect server has the following consequences:

- Hardware volume mappings cannot be created.
- Hardware volume mappings cannot be changed.
- The hardware volume mappings cannot be automatically replicated to HSM for Windows® clients on all nodes of a cluster.
- The hardware volume mapping is not applied when you search for files at a remote IBM® Storage Protect server connection.

Mapping volumes

To continue HSM services when a volume or file server is renamed, you must create a hardware volume mapping.

About this task

Create a hardware volume mapping by using the following steps:

Procedure

1. In the HSM for Windows™ client GUI, click **Tools > Volume Mappings**.
The **Hardware Volume Mappings** window lists all local volumes that are assigned a drive letter and all MSCS cluster volumes that are online.
If remote IBM® Storage Protect server connections exist, you can see the hardware volume mappings that are defined on the remote HSM for Windows™ client. You can view the remote hardware volume mappings but not change them.
2. Select a volume and click **Create**.
3. Type the old host and volume information and click **OK**.
The new mapping is displayed in the **Hardware Volume Mapping Definition** window. The mapping applies to all nested volumes on the selected drive.
4. After defining all hardware volume mappings, click **Close**.
The **Reconfirmation** window displays all new mappings.
5. Test the mappings by clicking **Scan for Problems**.
This test checks for files on the IBM® Storage Protect server that is defined with the old mapping. The scan shows whether there are any migrated files at the old mapping. After the new mapping is applied, the migrated files at the old mapping are not accessible.

6. Click **Yes** to apply the changes.
All HSM services receive notifications and apply the new mappings. HSM commands apply the new mappings the next time the commands are started.

Displaying HSM listing files

You can filter and search HSM listing files by using the HSM for Windows™ client graphical user interface (GUI).

About this task

The HSM for Windows™ client records files that are processed by HSM operations. The records are saved in listing files. The HSM for Windows™ client GUI can filter the listing files and display the records that you want to see. You can arrange the output columns and sort by column. You can search by system identification (SID) number or by user information.

To display records from the listing files, complete the following steps.

Procedure

1. In the HSM for Windows™ client GUI, select **Tools > Search Listing Files**.
The menu displays a choice for the location of the listing files.
2. Choose the default location or select another location and browse to the directory of the listing files.
After you specify the location, a search window opens. The window contains tabs for HSM operations: migration, recall, retrieve, delete, move, and trace.
3. Click a tab for the HSM operation that you want to see.
To display records for all operations, select the **All** tab.
4. Set the filters to display the records that you want to see. Click **Search**.
A progress window displays the search status.
When you search for recall records, you can search for records that contain a user account that is recognized by the Windows system. First you must query whether the Windows system recognizes the user account.

If the Windows system recognizes the user account, you can use the user account to filter the HSM for Windows recall records. The HSM for Windows recall records for that user contain the user account, but do not contain the SID.

When the search process is complete, the search results are displayed.

What to do next

If the search yields too many records, you can search again and specify more restrictive filters. You can hide and arrange the columns in the results window. You can order the records in a column. You can save the results to a file.

HSM for Windows™ client REST API and web client

The HSM for Windows™ client uses the REST API, which is used to access metrics through a web application, and the web client, which is used to supervise all the HSM-managed file servers.

Monitoring HSM-managed volumes with Prometheus and Grafana open platforms

You can view metrics for volumes that are managed by HSM for Windows™. The metrics can help you determine whether enough free space is available and whether more files can be migrated to release space.

Before you begin

OpenSSL software Version 1.1.1 must be installed on all HSM-managed file servers.

About this task

The HSM for Windows™ client provides REST API functionality to access metrics through a web application. The `hsmnet.exe` service provides recall and migration counts and volume space data. For example, the Prometheus open source system can connect to HSM-managed file servers to monitor recall and migration activity, and free space. You can display the HSM for Windows™ metrics in the Grafana Analytics Platform by connecting to the Prometheus server.

The following metrics are available:

- `hsm_backup_bytes`. The number of bytes that were backed up before migration by HSM for Windows™ to the backup storage pool.
- `hsm_backup_files`. The number of files that were backed up before migration by HSM for Windows™ to the backup storage pool.
- `hsm_moved_backend_bytes`. The number of bytes of backend files that were moved within the same IBM® Storage Protect server or different IBM® Storage Protect servers.
- `hsm_moved_backend_files`. The number of backend files that were moved within the same IBM® Storage Protect server or different IBM® Storage Protect servers.
- `hsm_moved_stubs`. The number of stub files that were moved to another volume or file server.
- `hsm_migrated_bytes`. The number of bytes that were migrated by HSM for Windows™ to the HSM archive.
- `hsm_migrated_bytes_details`. The details of the bytes that were migrated by HSM for Windows™ to the HSM archive.
- `hsm_migrated_files`. The number of files that were migrated by HSM for Windows™ to the HSM archive.
- `hsm_migrated_files_details`. The details of the files that were migrated by HSM for Windows™ to the HSM archive.
- `hsm_network_bytes`. The number of bytes sent through the network between the HSM for Windows™ client and the IBM® Storage Protect.
- `hsm_recalled_bytes`. The number of bytes that were recalled by HSM for Windows™ from the HSM archive.
- `hsm_recalled_files`. The number of files that were recalled by HSM for Windows™ from the HSM archive.
- `hsm_retrieved_bytes`. The number of bytes that were retrieved by HSM for Windows™ from the HSM archive.
- `hsm_retrieved_files`. The number of files that were retrieved by HSM for Windows™ from the HSM archive.
- `hsm_volume_free_bytes`. The amount of free space in bytes on the HSM for Windows™ volume.
- `hsm_volume_used_bytes`. The amount of used space in bytes on the HSM for Windows™ volume.
- `hsm_volume_used_percent`. The percentage of space that is used on the HSM for Windows™ volume.

The metrics can be displayed in the Prometheus Time Series Collection and its processing server, and in a Grafana dashboard that is configured for HSM for Windows™ metrics.

Monitoring HSM-managed volumes is helpful when HSM for Windows™ is installed on a new file server, and you want to optimize migration processes. The monitoring is also helpful when file systems have been managed by HSM for Windows™ for several years, but the migration processes are not releasing enough space.

Procedure

1. To enable secure communication between the web server system and HSM-managed file servers, create certificates with a key size of 2048 bits by taking the following actions:
 - a. Locate the `dsmcertificate.exe` tool, which is installed with the HSM for Windows™ client in the `hsmclient` installation folder.
 - b. On each node, run the `dsmcertificate.exe` tool and complete the steps as described in the readme file in the following directory: `installation_path/hsmclient/networking/README.html`, where `installation_path` is the path where HSM for Windows™ is installed.
2. Start the `hsmnet` service on each HSM for Windows™ node.

Tip: By default, the `hsmnet` service is not configured to start automatically during startup.

3. On the web server system, update the Prometheus configuration file with the configuration data for HSM for Windows™ as described in the readme file `installation_path/hsmclient/networking/README.html`.
4. Restart the Prometheus server.
5. Open the Prometheus Time Series Collection and processing server in a browser, and add graphs with HSM metrics. If HSM metrics are not listed, the connection might not be established with the processing server.
6. If the Grafana Analytics Platform is installed, you can define queries to the Prometheus time-series database, and select the metrics to be included in graphs. An example dashboard, `HSM-grafana-dashboard.json` is available in the folder `installation_path/hsmclient/networking`, and can be imported from the Grafana web interface.

Result

By viewing the metrics in trend charts and other graphs, you can answer the following questions:

- Whether enough free space is available on volumes.
- Whether more files can be migrated to release space
- Whether the number of recalled files is increasing because more files were migrated
- Whether additional space is required on volumes based on a limit defined by your environment
- Whether files are backed up before migration
- What network traffic is caused by the HSM for Windows™ client
- Whether a stub moving process to another file server or volume has finished.

Configuring the HSM for Windows™ client web GUI

Administrators can use a web browser to log in to the HSM for Windows™ client web GUI, access one of the HSM nodes in the company's installation, and supervise all the HSM-managed file servers.

Before you begin

OpenSSL software version 3.0 or later must be installed on all HSM-managed file servers.

About this task

The HSM for Windows™ client web GUI is an optional, installable component that comes with the HSM for Windows™ client installer package. The HSM for Windows™ client web GUI is installed on one HSM node that acts as the proxy-server node. All other HSM nodes can act as proxy-client nodes and require registration at the proxy-server node.

The HSM for Windows™ client web GUI consists of:

- The **Login** screen.
- The **Dashboard** with overview information of the HSM nodes in the company's installation.
- The **Activity** page with live information per file server volume in the **System Information** tab, and HSM information in the **Volume Scans** tab.
- The **Alerts** page with administration log information for all file servers, ordered in a timely fashion.
- The **Configuration** page for installation overview, and threshold migration information.

For more information, see ["Getting started with HSM for Windows client web GUI" on page 18.](#)

Procedure

- To configure the HSM for Windows™ client web GUI, see the instructions as described in the readme file *installation_path/hsmclient/networking/README.html*.

HSM for Windows® commands

The HSM for Windows™ client has several commands that you can run from a **Command Prompt** window. With these commands, you can do most of the tasks that you can do with the GUI.

“HSM for Windows™ client Command Prompt window commands” on page 75 summarizes the HSM commands.

| Table 10: HSM for Windows™ client Command Prompt window commands | |
|---|---|
| Command | Description |
| dsmclc.exe | Use this command to run a migration job from the Command Prompt window. You can also list files and file spaces, and set the level of information that is saved in log, trace, and list files. |
| dsminfo.exe | Use this command to list various settings of your installation. List the version of libraries, actual log level settings, the operating system version, and disk information. |
| dsmfileinfo.exe | Use this command to list attributes of migrated and non-migrated files. |
| dsmfind.exe | Use this command to list files that are eligible by a job file or that correspond to a pattern. |
| dsmhsmclc.exe | Use this command to manage reconciliation and threshold migration. You can also set the level of information that is saved in log, trace, and list files. |
| dsmmove.exe | Use this command to move stub files to another location. If the other location is managed by a different IBM Storage Protect™ server, the migrated file data is moved to the new IBM Storage Protect™ server. |
| dsmquota.exe | Use this command to display user and group quotas or to reset the quota recall counter for one or more users. |
| dsmtool.exe | Use this command to display the quantity, size, and expiration period of migrated objects in IBM Storage Protect™ storage. |

You can complete additional tasks without using the HSM for Windows GUI by manually editing job and configuration files. For more information about using the command line interface, see [technote 1381502](#).

Entering command parameters

Case sensitivity

Command options are not case-sensitive. You can type them in uppercase or lowercase.

Minimum abbreviation

In the syntax diagrams, the minimum abbreviation of a command option is printed in uppercase. For example, if the syntax diagram includes the option **-UNCONFIGUREReconcile**, the minimum abbreviation is UNCONFIGURER.

Restriction for running a command again

Wait for a command to finish before you enter that command again. If you enter a command when an instance of that command is running, you can get the following error message:

```
Could not open log file.  
Exiting.
```

Command parameters override default and job settings

The parameter values that you enter with a command override the values that you set in a job file or by using the configuration wizard.

Client return codes

The HSM for Windows™ client command-line interface exits with return codes that reflect the success or failure of the operation.

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

Return codes have the following meanings:

| Table 11: An explanation of client return codes | |
|---|--|
| Code | Explanation |
| 0 | All operations completed successfully. |
| 4 | The operation completed successfully, but some files were not processed. There were no other errors or warnings. This return code is common. In most cases, files are not processed for the following reasons: <ul style="list-style-type: none">• The file satisfies an entry in an exclude list. Excluded files generate log entries only during selective backups.• The file was in use by another application and could not be accessed by the client.• The file changed during the operation to an extent prohibited by the copy serialization attribute. |
| 8 | The operation completed with at least one warning message. Review the <code>dsmeerror.log</code> file to determine what warning messages were issued and to assess their effect on the operation. |
| 12 | The operation completed with at least one error message (except for error messages for skipped files). For scheduled events, the status is Failed. Review the <code>dsmeerror.log</code> file to determine what error messages were issued and to assess their effect on the operation. Generally, this return code means that the error was severe enough to prevent the successful completion of the operation. For example, an error that prevents an entire file system from being processed yields return code 12. |

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

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

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

dsmc1c.exe

The **dsmc1c.exe** command starts a migration job or a list migration, recalls and retrieves selected migrated files, creates and lists file spaces, lists and deletes migrated files, lists management classes, and creates server connections.

The optional parameters can be entered in any order.

Display help for the command by using the **help** parameter:

```
dsmc1c help
```

dsmc1c createfilespace

The **dsmc1c.exe** command with the **createfilespace** parameter creates a new file space on an IBM® Storage Protect server. After you create a file space, you can migrate files to that file space.

Syntax

► DSMCLC.exe — CREATEFILESPEC — -g — *new_filespace* — — -l — *loglevel* ►

Parameters

-g *new_filespace*

Specify a new file space name on IBM® Storage Protect storage.

-L *loglevel*

Specify the type of information that is to be recorded in logs and trace files. You can specify one or more values with no commas or blank space separators. Severe and error messages are always recorded. The default combination is severe, error, warning, information, and library (SEWIL). The following values are valid:

- C (event)
- D (debug)
- E (error)
- F (flush)
- I (information)
- K (driver)
- L (library)
- S (severe)
- T (trace)
- U (user)
- W (warning)
- X (dump)

Examples

Task

Create a file space: def-hsm02.

Command: dsmc1c createfilespace -g def-hsm02

Task

Display help for the **dsmc1c.exe** command.

Command: dsmc1c help

Task

Change the information that is recorded in log and trace files to the default.

Command: dsmc1c -l

dsmc1c defaults

The **dsmc1c.exe** command with the **defaults** parameter displays the default values for **dsmc1c.exe** command options.

Syntax

► DSMCLC.exe — DEFAULTS — — -L — *loglevel* ►

Parameters

-L *loglevel*

Specify the type of information that is to be recorded in logs and trace files. You can specify one or more values with no commas or blank space separators. Severe and error messages are always recorded. The default combination is severe, error, warning, information, and library (SEWIL). The following values are valid:

C (event)
D (debug)
E (error)
F (flush)
I (information)
K (driver)
L (library)
S (severe)
T (trace)
U (user)
W (warning)
X (dump)

Examples

Task

Display the default values for **dsmc1c.exe** command options.

Command: dsmc1c defaults

Task

Display help for the **dsmc1c.exe** command.

Command: dsmc1c help

Task

Change the information that is recorded in log and trace files to the default.

Command: dsmc1c -l

dsmc1c delete

The **dsmc1c.exe** command with the **delete** parameter deletes migrated files from IBM® Storage Protect storage.

Syntax

➔ DSMCLC.exe — DELETE — -g — *filespace* — *search_pattern* — *connection_options* — -l — *loglevel* ➔

Parameters

-g *filespace*

Specify a file space on IBM® Storage Protect storage. The file space name is case sensitive.

search_pattern

Specify a pattern for migrated objects. All migrated objects that match the pattern are included in the operation. There are several parts to a search pattern. Some parts are required; some parts are optional. Separate the parts with a blank space. Search pattern elements are case-sensitive. If there is no hardware mapping, you can use wildcard characters asterisk (*) and question mark (?).

volume_pattern

Specify a pattern that matches volume names. The volume pattern is required. If the volume pattern contains blank spaces, enclose the pattern with quotation marks.

If there is a hardware mapping for the volume, you must specify the file-server host name and drive letter without wildcard characters.

directory_pattern

Specify a pattern that matches directory names. The directory pattern is required. If the directory pattern contains blank spaces, enclose the pattern with quotation marks.

file_pattern

Specify a pattern that matches file names. The file pattern is optional. If the volume pattern contains blank spaces, enclose the pattern with quotation marks.

-version *number*

Specify a file version.

The **version** parameter is optional. If you do not specify a version, all versions are deleted.

connection_options

If the operation involves a remote file server, you must specify an IBM® Storage Protect connection.

You can specify a connection by specifying the two parts of a connection pair or by specifying a shortcut.

Specify the two parts of a connection

Use the **h** and **u** parameters:

-h *TSM_host_name*

Specify the IBM® Storage Protect server part of a connection pair. The value of *TSM_host_name* is not case-sensitive. Specify *TSM_host_name* with the value of the **TCPSERVERADDRESS** option and the value of the **TCPPORT** option, separated with a colon. For example: 127.0.0.1:1500

-u *node_name*

Specify the IBM® Storage Protect node part of a connection pair. Use the same value that you used to define the IBM® Storage Protect server connection. If the connection to the IBM® Storage Protect server was configured with the **asnodename** option, specify the value of the **asnodename** option. If the connection was configured without the **asnodename** option, specify the value of the **nodename** option. The value of *node_name* is not case-sensitive.

Specify a connection shortcut

Instead of specifying the host name and node name parts of a connection, you can specify a connection shortcut. Use the **c** parameter to specify a connection shortcut:

-c *shortcut*

The *shortcut* value is one or two characters and is generated by the HSM for Windows™ client. Connection shortcuts include these examples:

- *l* (local)
- *r1* (remote connection 1)
- *r2* (remote connection 2)

Tip: Run **dsmc1c help** to display connection shortcuts.

-L loglevel

Specify the type of information that is to be recorded in logs and trace files. You can specify one or more values with no commas or blank space separators. Severe and error messages are always recorded. The default combination is severe, error, warning, information, and library (SEWIL). The following values are valid:

C (event)
D (debug)
E (error)
F (flush)
I (information)
K (driver)
L (library)
S (severe)
T (trace)
U (user)
W (warning)
X (dump)

Examples

Task

Delete the migrated objects in the `c:\projects\2005\` directory. The migrated objects are in file space `def-hsm01`.

Command: `dsmc1c delete -g def-hsm00 c: \projects\2005`

Task

Delete all migrated *.doc files in the `\projects\2011\` directory of a remote IBM® Storage Protect server. The remote IBM® Storage Protect is identified by connection shortcut `r2`. The migrated files are in file space `def-hsm01`.

Command: dsmc1c delete -c r2 -g def-hsm01 \\remote_file_server\G\$ \projects\2011\ *.doc

Task

Display help for the **dsmc1c.exe** command.

Command: dsmc1c help

Task

Change the information that is recorded in log and trace files to the default.

Command: dsmc1c -l

dsmc1c legend

The **dsmc1c.exe** command with the **legend** parameter displays legends for table headers. Some table header text is abbreviated; the legends explain the table headers.

The **legend** parameter displays legends for tables that are output from **dsmc1c.exe** commands.

Syntax

►— DSMCLC.exe — LEGEND —◄

Examples

Task

Display legends for tables that are output from **dsmc1c.exe** commands.

Command: dsmc1c legend

Result:

```
Table column headers by command:

list, retrieve
  SIZE      file size in KB
  V          current file version
  S          file security availability
  FILENAME  file name

migrate, migratelist, recall, recalllist
  SIZE      file size in KB
  V          migrated file version
  FILENAME  file name

listfilespace
  NAME      file space name
  OCCUPANCY file space occupancy

listmgmtclasses
  NAME      management class name
  POLICY    management class policy
```

Task

Display help for the **dsmc1c.exe** command.

Command: dsmc1c help

dsmc1c list

The **dsmc1c.exe** command with the **list** parameter lists files that were migrated to IBM® Storage Protect storage.

For each migrated file, the following information is displayed:

- File size
- (V) File version number
- (S) Whether security attributes were migrated. A plus sign (+) indicates that security attributes were migrated.

- (D) Whether Windows™ alternate data stream data was migrated. A plus sign (+) indicates that Windows™ alternate data stream data was migrated.
- File path

Syntax

► **DSMCLC.exe** — **LIST** — **-g** — *filesystem* — *search_pattern* — *connection_options* — **-l** — *loglevel* — **-v** — ►

Parameters

connection_options

If the operation involves a remote file server, you must specify an IBM® Storage Protect connection.

You can specify a connection by specifying the two parts of a connection pair or by specifying a shortcut.

Specify the two parts of a connection

Use the **h** and **u** parameters:

-h *TSM_host_name*

Specify the IBM® Storage Protect server part of a connection pair. The value of *TSM_host_name* is not case-sensitive. Specify *TSM_host_name* with the value of the **TCPSERVERADDRESS** option and the value of the **TCPPORT** option, separated with a colon. For example: 127.0.0.1:1500

-u *node_name*

Specify the IBM® Storage Protect node part of a connection pair. Use the same value that you used to define the IBM® Storage Protect server connection. If the connection to the IBM® Storage Protect server was configured with the **asnodename** option, specify the value of the **asnodename** option. If the connection was configured without the **asnodename** option, specify the value of the **nodename** option. The value of *node_name* is not case-sensitive.

Specify a connection shortcut

Instead of specifying the host name and node name parts of a connection, you can specify a connection shortcut. Use the **c** parameter to specify a connection shortcut:

-c *shortcut*

The *shortcut* value is one or two characters and is generated by the HSM for Windows™ client. Connection shortcuts include these examples:

- *l* (local)
- *r1* (remote connection 1)
- *r2* (remote connection 2)

Tip: Run **dsmc1c help** to display connection shortcuts.

-g *filesystem*

Specify a file space on IBM® Storage Protect storage. The file space name is case sensitive.

search_pattern

Specify a pattern for migrated objects. All migrated objects that match the pattern are included in the operation. There are several parts to a search pattern. Some parts are required; some parts are optional. Separate the parts with a blank space. Search pattern elements are case-sensitive. If there is no hardware mapping, you can use wildcard characters asterisk (*) and question mark (?).

volume_pattern

Specify a pattern that matches volume names. The volume pattern is required. If the volume pattern contains blank spaces, enclose the pattern with quotation marks.

If there is a hardware mapping for the volume, you must specify the file-server host name and drive letter without wildcard characters.

directory_pattern

Specify a pattern that matches directory names. The directory pattern is required. If the directory pattern contains blank spaces, enclose the pattern with quotation marks.

file_pattern

Specify a pattern that matches file names. The file pattern is optional. If the volume pattern contains blank spaces, enclose the pattern with quotation marks.

-version number

Specify a file version.

The **version** parameter is optional. If you do not specify a version, all versions are listed.

-L loglevel

Specify the type of information that is to be recorded in logs and trace files. You can specify one or more values with no commas or blank space separators. Severe and error messages are always recorded. The default combination is severe, error, warning, information, and library (SEWIL). The following values are valid:

- C (event)
- D (debug)
- E (error)
- F (flush)
- I (information)
- K (driver)
- L (library)
- S (severe)
- T (trace)
- U (user)
- W (warning)
- X (dump)

-v

Display verbose output.

Examples

Task

List all *.doc migrated files in the c:\big projects\2009\ directory. The migrated files are in file space def-hsm01.

Command: dsmc1c list -g def-hsm01 c: "\big projects\2009" *.doc

Task

List all migrated *.doc files in the \projects\2011\ directory of a remote IBM® Storage Protect server. The remote IBM® Storage Protect is identified by connection shortcut r2. The migrated files are in file space def-hsm01.

Command: dsmc1c list -c r2 -g def-hsm01 \\remote_file_server\G: \projects\2011\ *.doc

Task

Display help for the **dsmc1c.exe** command.

Command: dsmc1c help

Task

Change the information that is recorded in log and trace files to the default.

Command: dsmc1c -l

dsmc1c listfilespace

The **dsmc1c.exe** command with the **listfilespace** parameter lists file spaces on an IBM® Storage Protect server. The HSM for Windows™ client lists all file spaces that you are authorized to see. The command indicates the occupancy of the file space.

The occupancy data that is displayed by the command **dsmc1c listfilespace** is the sum of file sizes of all migrated files for a file space. Occupancy also includes information for managing the migrated files.

Compression, data deduplication, and expirations on the IBM® Storage Protect server are not reflected in the statistics from the **dsmc1c listfilespace** command. The occupancy data is refreshed when you run the **dsmtool** command with the **occupancy** or **statistic** parameters.

For details on expirations due to copy group settings, see [Technote 1330160](#).

Syntax

► DSMCLC.exe — LISTFILESPPACES — *file_space_pattern* — *connection_options* — -l — *loglevel* ►

Parameters

connection_options

If the operation involves a remote file server, you must specify an IBM® Storage Protect connection.

You can specify a connection by specifying the two parts of a connection pair or by specifying a shortcut.

Specify the two parts of a connection

Use the **h** and **u** parameters:

-h *TSM_host_name*

Specify the IBM® Storage Protect server part of a connection pair. The value of *TSM_host_name* is not case-sensitive. Specify *TSM_host_name* with the value of the **TCPSERVERADDRESS** option and the value of the **TCPPORT** option, separated with a colon. For example: 127.0.0.1:1500

-u *node_name*

Specify the IBM® Storage Protect node part of a connection pair. Use the same value that you used to define the IBM® Storage Protect server connection. If the connection to the IBM® Storage Protect server was configured with the **asnodename** option, specify the value of the **asnodename** option. If the connection was configured without the **asnodename** option, specify the value of the **nodename** option. The value of *node_name* is not case-sensitive.

Specify a connection shortcut

Instead of specifying the host name and node name parts of a connection, you can specify a connection shortcut. Use the **c** parameter to specify a connection shortcut:

-c *shortcut*

The *shortcut* value is one or two characters and is generated by the HSM for Windows™ client. Connection shortcuts include these examples:

- *l* (local)
- *r1* (remote connection 1)
- *r2* (remote connection 2)

Tip: Run **dsmc1c help** to display connection shortcuts.

file_space_pattern

Specify a pattern for file spaces. If there is a blank space in the pattern, surround the pattern with quotation marks. Search pattern elements are case-sensitive. You can use wildcard characters * and ?.

-L loglevel

Specify the type of information that is to be recorded in logs and trace files. You can specify one or more values with no commas or blank space separators. Severe and error messages are always recorded. The default combination is severe, error, warning, information, and library (SEWIL). The following values are valid:

C (event)
D (debug)
E (error)
F (flush)
I (information)
K (driver)
L (library)
S (severe)
T (trace)
U (user)
W (warning)
X (dump)

Examples

Task

List all file spaces that you are authorized to see.

Command: `dsmc1c listfilespace`

Task

Display help for the **dsmc1c.exe** command.

Command: `dsmc1c help`

Task


Change the information that is recorded in log and trace files to the default.

Command: `dsmc1c -l`

dsmc1c listmgmtclasses

The **dsmc1c.exe** command with the **listmgmtclasses** parameter lists management classes that contain an archive copy group. (A management class must contain an archive copy group to store files that are migrated.) You can use a pattern to filter management class names.

Syntax



Parameters

mgmt_class_pattern

Specify a pattern for management classes. If there is a blank space in the pattern, surround the pattern with quotation marks. Search pattern elements are case-sensitive. You can use wildcard characters * and ?.

connection_options

If the operation involves a remote file server, you must specify an IBM® Storage Protect connection.

You can specify a connection by specifying the two parts of a connection pair or by specifying a shortcut.

Specify the two parts of a connection

Use the **h** and **u** parameters:

-h *TSM_host_name*

Specify the IBM® Storage Protect server part of a connection pair. The value of *TSM_host_name* is not case-sensitive. Specify *TSM_host_name* with the value of the **TCPSERVERADDRESS** option and the value of the **TCPPORT** option, separated with a colon. For example: `127.0.0.1:1500`

-u *node_name*

Specify the IBM® Storage Protect node part of a connection pair. Use the same value that you used to define the IBM® Storage Protect server connection. If the connection to the IBM® Storage Protect server was configured with the **asnodename** option, specify the value of the **asnodename** option. If the connection was configured without the **asnodename** option, specify the value of the **nodename** option. The value of *node_name* is not case-sensitive.

Specify a connection shortcut

Instead of specifying the host name and node name parts of a connection, you can specify a connection shortcut. Use the **c** parameter to specify a connection shortcut:

-c *shortcut*

The *shortcut* value is one or two characters and is generated by the HSM for Windows™ client. Connection shortcuts include these examples:

- `l` (local)

- *r1* (remote connection 1)
- *r2* (remote connection 2)

Tip: Run **dsmc1c help** to display connection shortcuts.

-L loglevel

Specify the type of information that is to be recorded in logs and trace files. You can specify one or more values with no commas or blank space separators. Severe and error messages are always recorded. The default combination is severe, error, warning, information, and library (SEWIL). The following values are valid:

C (event)
D (debug)
E (error)
F (flush)
I (information)
K (driver)
L (library)
S (severe)
T (trace)
U (user)
W (warning)
X (dump)

-v

Display verbose output.

Examples

Task

List the properties of the DEFAULT management class.

Command: `dsmc1c listmgmtclasses DEFAULT`

Task

Display help for the **dsmc1c.exe** command.

Command: `dsmc1c help`

Task

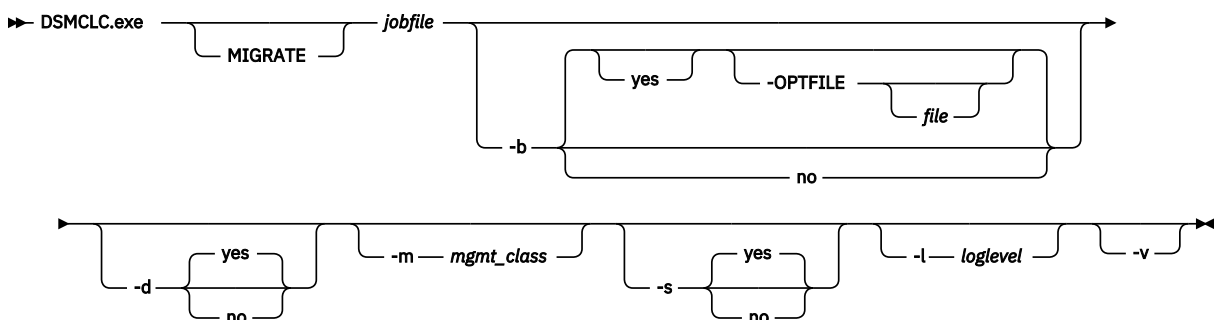
Change the information that is recorded in log and trace files to the default.

Command: `dsmc1c -l`

dsmc1c migrate

The **dsmc1c.exe** command with the **migrate** parameter starts a migration job. You can omit the **migrate** parameter, but you must specify the job file name.

Syntax



Parameters

jobfile

Specify a migration job file. You can specify a complete path, or only file name and file type, or only the file name. For example, the following commands specify the same job file:

- dsmclic c:\hsmclient\jobs\migrate011.osj
- dsmclic migrate011.osj
- dsmclic migrate011

-b yes|no

Specify whether files are backed up before migration. The default is the value that you set in the initial configuration wizard. If you use the option but do not specify yes or no, files are backed up before migration.

-d yes|no

Specify whether Windows™ alternate data stream (ADS) data is migrated when the file is migrated. The default is the value that you set in the initial configuration wizard. A yes value means that ADS data is migrated when the file is migrated. A no value means that ADS data is not migrated. If you use the option but do not specify yes or no, a yes value is assumed.

-OPTFILE file

Specify the path of an options file for backup before migration.
This option is valid only if you also specify backup before migration.

If *file* is not specified, the backup-archive client uses its default options file. This file value overrides the value that is configured in a migration job file.

-m mgmt_class

Specify a management class for the migration job or list migration. This value overrides the management class that is specified when the job was created. Specify DEFAULT to use the IBM Storage Protect™ server default management class of the active policy set.

-s yes|no

Specify whether file security attributes (ACL) are migrated when the file is migrated. The default is the value that you set in the initial configuration wizard. A yes value means that the ACL is migrated when the file is migrated. A no value means that the ACL is not migrated. If you use the option but do not specify yes or no, a yes value is assumed.

-L loglevel

Specify the type of information that is to be recorded in logs and trace files. You can specify one or more values with no commas or blank space separators. Severe and error messages are always recorded. The default combination is severe, error, warning, information, and library (SEWIL). The following values are valid:

C (event)
D (debug)
E (error)
F (flush)
I (information)
K (driver)
L (library)
S (severe)
T (trace)
U (user)
W (warning)
X (dump)

-v

Display verbose output.

Examples

Task

Migrate files with the job that is defined in c:\hsmclient\jobs\migrate011.osj.

Command: dsmclic c:\hsmclient\jobs\migrate011.osj

Task

Migrate files with the job that is defined in `c:\hsmclient\jobs\migrate011.osj`. Use management class MC2. The backup-archive client determine the options file, even if you specified another options file when you configured this job.

Command: `dsmc1c -m MC2 c:\hsmclient\jobs\migrate011.osj -optfile`

Task

Display help for the **dsmc1c.exe** command.

Command: `dsmc1c help`

Task

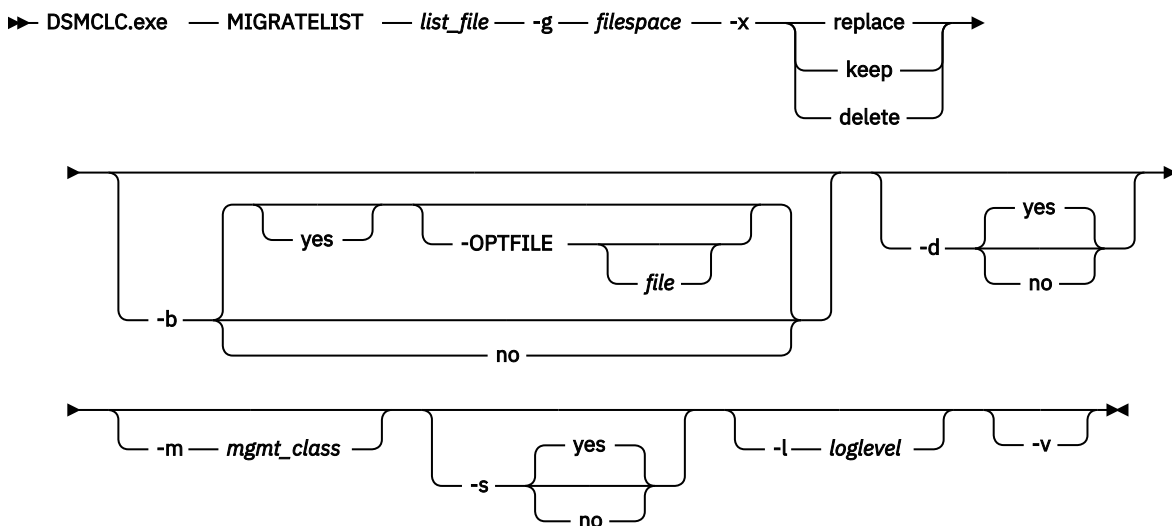
Change the information that is recorded in log and trace files to the default.

Command: `dsmc1c -l`

dsmc1c migratelist

The **dsmc1c.exe** command with the **migratelist** parameter migrates files that are listed in a list file.

Syntax



Parameters

list_file

Specify the path of a list file. The list file contains a list of files. Within the list, each file is on a separate line. Each file is identified by a complete path from the root. For example, `c:\projects\2009\budget1.xls`. The list file is not created by the HSM for Windows™ client GUI. The list can be encoded in ASCII or Unicode. If Unicode, the first 2 bytes must be the byte order mark (BOM).

-g filespace

Specify a file space on IBM® Storage Protect storage. The file space name is case sensitive.

-x

Specify an action on the file system after the file is migrated to IBM® Storage Protect storage:

REPLACE

Replace the migrated file with a stub file.

KEEP

Keep the complete file on the file system.

DELETE

Delete the file from the file system.

-b yes|no

Specify whether files are backed up before migration. The default is the value that you set in the initial configuration wizard. If you use the option but do not specify `yes` or `no`, files are backed up before migration.

-d yes|no

Specify whether Windows™ alternate data stream (ADS) data is migrated when the file is migrated. The default is the value that you set in the initial configuration wizard. `Ayes` value means that ADS data is migrated when the file is migrated. `Ano` value means that ADS data is not migrated. If you use the option but do not specify `yes` or `no`, `ayes` value is assumed.

-OPTFILE file

Specify the path of an options file for backup before migration.
This option is valid only if you also specify backup before migration.

If *file* is not specified, the backup-archive client uses its default options file. This file value overrides the value that is configured in a migration job file.

-m mgmt_class

Specify a management class for the migration job or list migration. This value overrides the management class that is specified when the job was created. Specify `DEFAULT` to use the IBM Storage Protect™ server default management class of the active policy set.

-s yes|no

Specify whether file security attributes (ACL) are migrated when the file is migrated. The default is the value that you set in the initial configuration wizard. `Ayes` value means that the ACL is migrated when the file is migrated. `Ano` value means that the ACL is not migrated. If you use the `-s` option but do not specify `yes` or `no`, the ACL is migrated when the file is migrated.

-L loglevel

Specify the type of information that is to be recorded in logs and trace files. You can specify one or more values with no commas or blank space separators. Severe and error messages are always recorded. The default combination is severe, error, warning, information, and library (SEWIL). The following values are valid:

- C (event)
- D (debug)
- E (error)
- F (flush)
- I (information)
- K (driver)
- L (library)
- S (severe)
- T (trace)
- U (user)
- W (warning)
- X (dump)

-v

Display verbose output.

Examples

Task

Migrate files in the list file `c:\hsmclient\jobs\xlsfiles.txt` to file space `def-hsm01`. Replace the migrated files with stubs. Back up files before migrating. Use options file `d:\backupAdmin\optionsFiles\backup_options_set3.opt`.

Command: `dsmc1c migratelist -g def-hsm01 -x replace c:\hsmclient\jobs\xlsfiles.txt -b -optfile d:\backupAdmin\optionsFiles\backup_options_set3.opt`

Task

Display help for the `dsmc1c.exe` command.

Command: `dsmc1c help`

Task

Change the information that is recorded in log and trace files to the default.

Command: `dsmc1c -l`


dsmc1c recall

The **dsmc1c.exe** command with the **recall** parameter recalls migrated files by searching for selected stub files on the file system.

The following restrictions apply:

- Hidden files are not recalled if you use wildcard characters to specify file names. You must specify the full path of hidden files.
- Stub files with the system attribute are not recalled. Files with the system attribute are not migrated.
- Windows™ alternate data stream (ADS) data is not recalled. The ADS data that is in the stub file is not changed when the primary data stream data is recalled.

Syntax



Parameters

find_pattern

Specify a stub-file path on the file system. All stub files that match the pattern are included in the operation. You can use wildcard characters asterisk (*) and question mark (?).

connection_options

Stub files that are in moving state point to a remote IBM® Storage Protect server. By default, files are recalled from the local IBM® Storage Protect server and any remote IBM® Storage Protect server that is indicated in any selected stub file. You can limit the recall operations to a single IBM® Storage Protect server by specifying a connection.

You can specify a connection by specifying the two parts of a connection pair or by specifying a shortcut.

Specify the two parts of a connection

Use the **h** and **u** parameters:

-h *TSM_host_name*

Specify the IBM® Storage Protect server part of a connection pair. The value of *TSM_host_name* is not case-sensitive. Specify *TSM_host_name* with the value of the **TCPSERVERADDRESS** option and the value of the **TCPPORT** option, separated with a colon. For example: 127.0.0.1:1500

-u *node_name*

Specify the IBM® Storage Protect node part of a connection pair. Use the same value that you used to define the IBM® Storage Protect server connection. If the connection to the IBM® Storage Protect server was configured with the **asnodename** option, specify the value of the **asnodename** option. If the connection was configured without the **asnodename** option, specify the value of the **nodename** option. The value of *node_name* is not case-sensitive.

Specify a connection shortcut

Instead of specifying the host name and node name parts of a connection, you can specify a connection shortcut. Use the **c** parameter to specify a connection shortcut:

-c *shortcut*

The *shortcut* value is one or two characters and is generated by the HSM for Windows™ client. Connection shortcuts include these examples:

- *l* (local)
- *r1* (remote connection 1)
- *r2* (remote connection 2)

Tip: Run **dsmc1c help** to display connection shortcuts.

-n stub_count

Specify the number of stub files that are processed in a single recall block. The stubs are sorted to optimize recall from tape devices. The IBM Storage Protect™ server locks a sequential storage device while the files in the recall block are recalled. A smaller value of *stub_count* allows other applications more frequent opportunities to access the device.

The default value is 5000.

A value of 0 specifies an unlimited block size. The IBM Storage Protect™ server locks a sequential storage device until all migrated files in the list file are recalled.

-r

Recurse into subdirectories to search for matching file names.

-t

Test the recall for space requirements. Files are not recalled. The HSM for Windows client calculates the space that is required to recall the files and identifies orphan files. Migration candidates are displayed.

-L loglevel

Specify the type of information that is to be recorded in logs and trace files. You can specify one or more values with no commas or blank space separators. Severe and error messages are always recorded. The default combination is severe, error, warning, information, and library (SEWIL). The following values are valid:

- C (event)
- D (debug)
- E (error)
- F (flush)
- I (information)
- K (driver)
- L (library)
- S (severe)
- T (trace)
- U (user)
- W (warning)
- X (dump)

Examples

Task

Calculate how much disk space is required to recall all migrated *.xls files in c:\projects\2013\accounting\ and all subdirectories.

Command: dsmc1c recall c:\projects\2013\accounting*.xls -r -t

Task

Recall all migrated *.xls files in c:\projects\2013\accounting\ and all subdirectories. Limit recalls to 500 files per block.

Command: dsmc1c recall c:\projects\2013\accounting*.xls -r -n 500

Task

Display help for the **dsmc1c.exe** command.

Command: dsmc1c help

Task

Change the information that is recorded in log and trace files to the default.

Command: dsmc1c -l

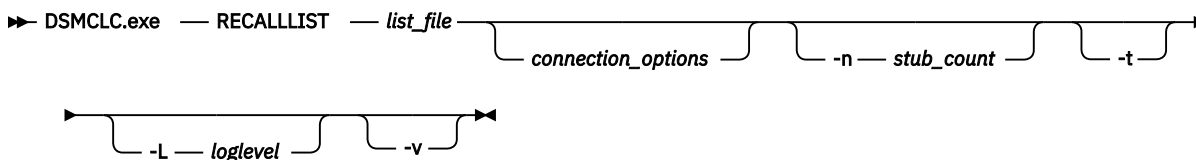
dsmc1c recalllist

The **dsmc1c.exe** command with the **recalllist** parameter recalls migrated files by searching the file system for files that are listed in a list file.

The following restrictions apply:

- Hidden files are not recalled if you use wildcard characters to specify file names. You must specify the full path of hidden files.
- Stub files with the system attribute are not recalled. Files with the system attribute are not migrated.
- Windows™ alternate data stream (ADS) data is not recalled. The ADS data that is in the stub file is not changed when the primary data stream data is recalled.

Syntax



Parameters

list file

Specify the path of a list file. The list file contains a list of files. Within the list, each file is on a separate line. Each file is identified by a complete path from the root. For example, c:\projects\2009\budget1.xls. The list file is not created by the HSM for Windows™ client GUI. The list can be encoded in ASCII or Unicode. If Unicode, the first 2 bytes must be the byte order mark (BOM).

connection options

Stub files that are in moving state point to a remote IBM® Storage Protect server. By default, files are recalled from the local IBM® Storage Protect server and any remote IBM® Storage Protect server that is indicated in any selected stub file. You can limit the recall operations to a single IBM® Storage Protect server by specifying a connection.

You can specify a connection by specifying the two parts of a connection pair or by specifying a shortcut.

Specify the two parts of a connection

Use the **h** and **u** parameters:

-h *TSM host name*

Specify the IBM® Storage Protect server part of a connection pair. The value of *TSM_host_name* is not case-sensitive. Specify *TSM_host_name* with the value of the **TCPSERVERADDRESS** option and the value of the **TCPPORT** option, separated with a colon. For example: 127.0.0.1:1500

-u node name

Specify the IBM® Storage Protect node part of a connection pair. Use the same value that you used to define the IBM® Storage Protect server connection. If the connection to the IBM® Storage Protect server was configured with the **asnodename** option, specify the value of the **asnodename** option. If the connection was configured without the **asnodename** option, specify the value of the **nodename** option. The value of *node name* is not case-sensitive.

Specify a connection shortcut

Instead of specifying the host name and node name parts of a connection, you can specify a connection shortcut. Use the `c` parameter to specify a connection shortcut:

-c shortcut

The *shortcut* value is one or two characters and is generated by the HSM for Windows™ client. Connection shortcuts include these examples:

- l (local)
- $r1$ (remote connection 1)
- $r2$ (remote connection 2)

Tip: Run `dsmc1c help` to display connection shortcuts.

-n stub_count

Specify the number of stub files that are processed in a single recall block. The stubs are sorted to optimize recall from tape devices. The IBM Storage Protect™ server locks a sequential storage device while the files in the recall block are recalled. A smaller value of *stub_count* allows other applications more frequent opportunities to access the device.

The default value is 5000.

A value of 0 specifies an unlimited block size. The IBM Storage Protect™ server locks a sequential storage device until all migrated files in the list file are recalled.

-t

Test the recall for space requirements. Files are not recalled. The HSM for Windows client calculates the space that is required to recall the files and identifies orphan files. Migration candidates are displayed.

-L loglevel

Specify the type of information that is to be recorded in logs and trace files. You can specify one or more values with no commas or blank space separators. Severe and error messages are always recorded. The default combination is severe, error, warning, information, and library (SEWIL). The following values are valid:

- C (event)
- D (debug)
- E (error)
- F (flush)
- I (information)
- K (driver)
- L (library)
- S (severe)
- T (trace)
- U (user)
- W (warning)
- X (dump)

-v

Display verbose output.

Examples

Task

Calculate how much disk space is required to recall all migrated files that are listed in `c:\lists\stub-files-for-recall.lst`.

Command: `dsmc1c recalllist -t c:\lists\stub-files-for-recall.lst`

Task

Recall all migrated files that are listed in `c:\lists\stub-files-for-recall.lst`. Limit recalls to 500 files per block.

Command: `dsmc1c recalllist c:\lists\stub-files-for-recall.lst -n 500`

Task

Display help for the **dsmc1c.exe** command.

Command: `dsmc1c help`

Task

Change the information that is recorded in log and trace files to the default.

Command: `dsmc1c -l`

dsmc1c register

The **dsmc1c.exe** command with the **register** parameter creates an IBM® Storage Protect server connection or sets the password of an existing connection.

Before you run the **dsmc1c register** command to create a connection, you must create an options file for the connection. To create a connection to the default IBM® Storage Protect server, create the options file in the HSM for Windows™ client installation directory. The name of this option file must be `dsm.opt`. To create a remote connection, create the options file in the `\config\` subdirectory of the HSM for Windows™ client installation directory. The options file must be file type `.opt`. The options file must contain the following options:

- **tcpserveraddress** *address*
- **tcpport** *port*

- **commethod** *tcpip*
- **passwordaccess** *generator passwordaccessprompt*
- **nodename** *name*.
The node must be registered with the IBM® Storage Protect server.

An optional entry in the options file is **asnodename** *name*. This node must be granted proxy authority.

Tip: To display the connections that are defined in the options files, enter the **dsmc1c** command with no parameters.

Syntax

```
➔ DSMCLC.exe — REGISTER — -p — password ————— connection_options — -l — loglevel —➔
```

Parameters

-p password

Specify the password for the node.

connection_options

If the operation involves a remote file server, you must specify an IBM® Storage Protect connection.

You can specify a connection by specifying the two parts of a connection pair or by specifying a shortcut.

Specify the two parts of a connection

Use the **h** and **u** parameters:

-h *TSM_host_name*

Specify the IBM® Storage Protect server part of a connection pair. The value of *TSM_host_name* is not case-sensitive. Specify *TSM_host_name* with the value of the **TCPSERVERADDRESS** option and the value of the **TCPPORT** option, separated with a colon. For example: 127.0.0.1:1500

-u *node_name*

Specify the IBM® Storage Protect node part of a connection pair. Use the same value that you used to define the IBM® Storage Protect server connection. If the connection to the IBM® Storage Protect server was configured with the **asnodename** option, specify the value of the **asnodename** option. If the connection was configured without the **asnodename** option, specify the value of the **nodename** option. The value of *node_name* is not case-sensitive.

Specify a connection shortcut

Instead of specifying the host name and node name parts of a connection, you can specify a connection shortcut. Use the **c** parameter to specify a connection shortcut:

-c *shortcut*

The *shortcut* value is one or two characters and is generated by the HSM for Windows™ client. Connection shortcuts include these examples:

- *l* (local)
- *r1* (remote connection 1)
- *r2* (remote connection 2)

Tip: Run **dsmc1c help** to display connection shortcuts.

-L loglevel

Specify the type of information that is to be recorded in logs and trace files. You can specify one or more values with no commas or blank space separators. Severe and error messages are always recorded. The

default combination is severe, error, warning, information, and library (SEWIL). The following values are valid:

C (event)
D (debug)
E (error)
F (flush)
I (information)
K (driver)
L (library)
S (severe)
T (trace)
U (user)
W (warning)
X (dump)

Examples

Task

Create a connection to an IBM® Storage Protect server. You created an options file with the following values:

- **tcpserveraddress**HAMBURG_TSM
- **tcpport**1500
- **commmethod**tcpip
- **passwordaccess**generate
- **nodename**TSMNODE

Command: `dsmc1c register -h HAMBURG_TSM:1500 -u TSMNODE -p password`

Task

Display help for the **dsmc1c.exe** command.

Command: `dsmc1c help`

Task

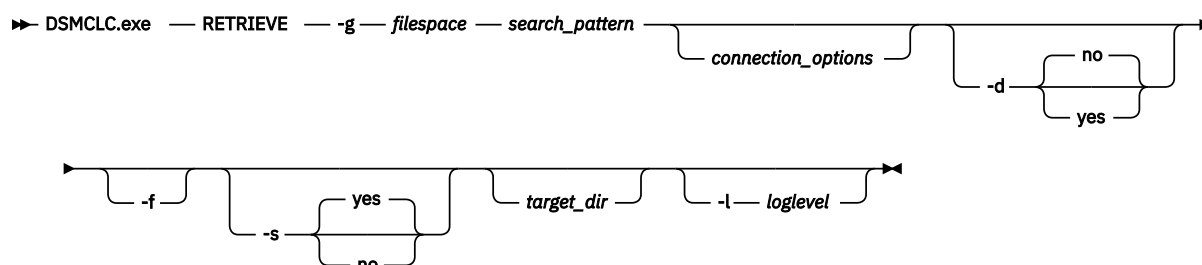
Change the information that is recorded in log and trace files to the default.

Command: `dsmc1c -l`

dsmc1c retrieve

The **dsmc1c.exe** command with the **retrieve** parameter retrieves migrated files from the file space of an IBM® Storage Protect server. Primary data stream (PDS) data and Windows™ alternate data stream (ADS) data is retrieved.

Syntax



Options

-g filesystem

Specify a file space on IBM® Storage Protect storage. The file space name is case sensitive.

search_pattern

Specify a pattern for migrated objects. All migrated objects that match the pattern are included in the operation. There are several parts to a search pattern. Some parts are required; some parts are optional. Separate the parts with a blank space. Search pattern elements are case-sensitive. If there is no hardware mapping, you can use wildcard characters asterisk (*) and question mark (?).

volume_pattern

Specify a pattern that matches volume names. The volume pattern is required. If the volume pattern contains blank spaces, enclose the pattern with quotation marks.

If there is a hardware mapping for the volume, you must specify the file-server host name and drive letter without wildcard characters.

directory_pattern

Specify a pattern that matches directory names. The directory pattern is required. If the directory pattern contains blank spaces, enclose the pattern with quotation marks.

file_pattern

Specify a pattern that matches file names. The file pattern is optional. If the volume pattern contains blank spaces, enclose the pattern with quotation marks.

-version number

Specify a file version.

The **version** parameter is optional. If you do not specify a version, only the most recent version is retrieved.

connection_options

If the operation involves a remote file server, you must specify an IBM® Storage Protect connection.

You can specify a connection by specifying the two parts of a connection pair or by specifying a shortcut.

Specify the two parts of a connection

Use the **h** and **u** parameters:

-h TSM_host_name

Specify the IBM® Storage Protect server part of a connection pair. The value of *TSM_host_name* is not case-sensitive. Specify *TSM_host_name* with the value of the **TCPSERVERADDRESS** option and the value of the **TCPPORT** option, separated with a colon. For example: 127.0.0.1:1500

-u node_name

Specify the IBM® Storage Protect node part of a connection pair. Use the same value that you used to define the IBM® Storage Protect server connection. If the connection to the IBM® Storage Protect server was configured with the **asnodename** option, specify the value of the **asnodename** option. If the connection was configured without the **asnodename** option, specify the value of the **nodename** option. The value of *node_name* is not case-sensitive.

Specify a connection shortcut

Instead of specifying the host name and node name parts of a connection, you can specify a connection shortcut. Use the **c** parameter to specify a connection shortcut:

-c shortcut

The *shortcut* value is one or two characters and is generated by the HSM for Windows™ client. Connection shortcuts include these examples:

- *l* (local)
- *r1* (remote connection 1)
- *r2* (remote connection 2)

Tip: Run **dsmc1c help** to display connection shortcuts.

-d no|yes

Specify whether Windows™ alternate data stream (ADS) data is retrieved when the file is retrieved. The default is the value that you set in the initial configuration wizard. **Ayesvalue** means that ADS data is retrieved when the file is retrieved. ADS data can be retrieved only if ADS data was migrated. **Ano**value means that ADS data is not retrieved. If you use the option but do not specify **yes** or **no**, **ayesvalue** is assumed.

-f

Force writing the retrieved file if a copy exists on the local volume.

If the stub file on the file system contains Windows™ alternate data stream (ADS) data, you must use the `f` option to retrieve the file. It is possible that the ADS data that is in the stub file is more recent than the ADS data that was migrated.

-s yes|no

Specify whether file security attributes (ACL) are migrated when the file is migrated. The default is the value that you set in the initial configuration wizard. A yes value means that the ACL is migrated when the file is migrated. A no value means that the ACL is not migrated. If you use the option but do not specify yes or no, a yes value is assumed.

target_dir

Specify a directory for the retrieved file. If you do not specify this option, the file is retrieved to the original path.

-L loglevel

Specify the type of information that is to be recorded in logs and trace files. You can specify one or more values with no commas or blank space separators. Severe and error messages are always recorded. The default combination is severe, error, warning, information, and library (SEWIL). The following values are valid:

- C (event)
- D (debug)
- E (error)
- F (flush)
- I (information)
- K (driver)
- L (library)
- S (severe)
- T (trace)
- U (user)
- W (warning)
- X (dump)

Examples

Task

Retrieve the migrated .xls files in the `c:\big projects\2009\` directory to a new path: `c:\projects\spreadsheets\`. The migrated copies are in file space `def-hsm01`.

Command: `dsmc1c retrieve -g def-hsm01 c: "\big projects\2009" *.xls c: \projects\spreadsheets.`

Spaces separate the three parts of the *search_pattern*: `c: "\big projects\2009" *.xls`. Because the *directory_pattern* (`\big projects\2009`) contains a blank space, it is enclosed in quotation marks.

Task

Display help for the `dsmc1c.exe` command.

Command: `dsmc1c help`

Task


Change the information that is recorded in log and trace files to the default.

Command: `dsmc1c -l`

dsmfileinfo.exe

Run the `dsmfileinfo.exe` program from a **Command Prompt** window to view file attributes.

Syntax

➤ `DSMFILEINFO.exe`  `file_path` ➤

Options

info_options

You can specify any of the following options. Separate options with a blank space.

| Table 12: Options for dsmfileinfo.exe | |
|---------------------------------------|---|
| Option | Description |
| -a | Display information for all options in this table |
| -d | Show alternate data streams |
| -i | Show file object ID |
| -ic | Create file object ID |
| -m | Calculate MD5 key (complete files only) |
| -q | Query backend version(s) (stub files only) |
| -r | Show reparse data (stub files only) |
| -rb | Show binary reparse data (stub files only) |
| -s | Show file security data |
| -sb | Show binary security data |
| -t | Show file times, size, and attributes (complete files only). This option is the default option. |

file_path

Specify the path of a complete file or a stub file. Specify only one file.

Examples

Task

Display the access time, creation time, modification time, size, and attributes of the file: c:\projects\2009\budget.xls.

Command: dsmfileinfo c:\projects\2009\budget.xls

Task

Create object ID for file c:\projects\2009\budget.xls.

Command: dsmfileinfo -ic c:\projects\2009\budget.xls

Task

Display binary security data for c:\projects\2009\budget.xls.

Command: dsmfileinfo -sb c:\projects\2009\budget.xls

Task

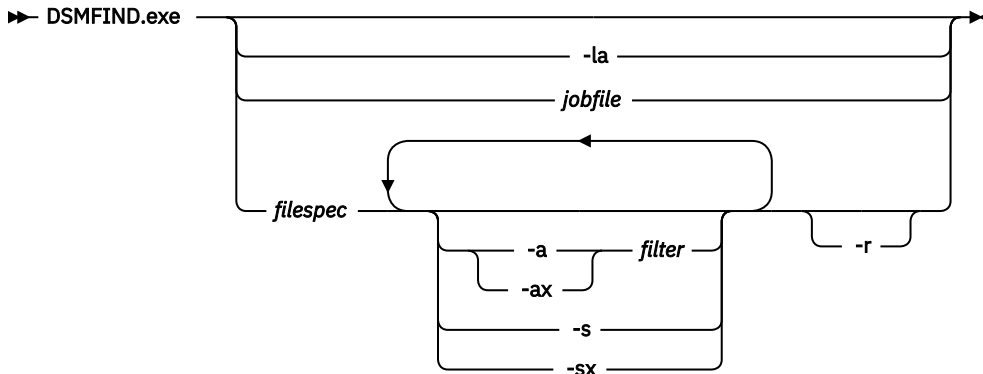
Display help for **dsmfileinfo.exe** command.

Command: dsmfileinfo

dsmfind.exe

Run the dsmfind.exe program from a **Command Prompt** window to show files that are described by a job file or by a file path and file attribute filters.

Syntax



Options

-la

List Windows® supported file attributes. Use the listed values to determine the filter for a **dsmfind.exe** command.

jobfile

Specify the path of a migration job file. The command displays all files that meet the criteria that is defined in a migration job file.

filter

Use this option with the attribute options (-a and -ax). Specify a filter for file attributes. The filter must be in the format 0xnnnnnnnn, where n is a hexadecimal number. You can combine file attributes. For example, the filter with value 0x00001600 is a combination of these file attributes:

- 0x00000200 (FILE_ATTRIBUTE_SPARSE_FILE)
- 0x00000400 (FILE_ATTRIBUTE_REPARSE_POINT)
- 0x00001000 (FILE_ATTRIBUTE_OFFLINE)

-aor-ai

Use this option with a filter. This option displays only files that have all attributes defined by the filter.

-ax

Use this option with a filter. This option excludes files that have all attributes defined by the filter.

-sor-si

This option displays stub files only. This is the same as -a 0x00001600. When stub files are created, stub files have these attributes:

- 0x00000200 (FILE_ATTRIBUTE_SPARSE_FILE)
- 0x00000400 (FILE_ATTRIBUTE_REPARSE_POINT)
- 0x00001000 (FILE_ATTRIBUTE_OFFLINE)

Note: Some anti-virus programs can remove the attribute FILE_ATTRIBUTE_OFFLINE from stub files.

-sx

This option excludes stub files. This is the same as -ax 0x00001600.

-r

The command displays files in all subdirectories.

Invoke the command with no options to display help for the command.

Examples

Task

Display all files that meet the criteria that is defined in the job file c : \hsmclient\jobs\migrate011.osj.

Command: dsmfind c:\hsmclient\jobs\migrate011.osj

Task

Display all Excel files in c:\projects\2009\.

Command: dsmfind c:\projects\2009*.xls

Task

Display all Excel files in c:\projects\ and all subdirectories.

Command: dsmfind c:\projects*.xls -r

Task

Display all stub files in c:\projects\ and all subdirectories.

Command: dsmfind c:\projects\ -r -s

Task

Display all read-only stub files in c:\projects\ and all subdirectories. Read-only files have attribute FILE_ATTRIBUTE_READONLY (0x00000001). Read-only stub files with other attributes are not displayed. Only files with combined attributes of 0x00001601 are displayed.

Command: dsmfind c:\projects\ -r -s -a 0x00000001

Task

Display help for the **dsmfind.exe** command.

Command: dsmfind

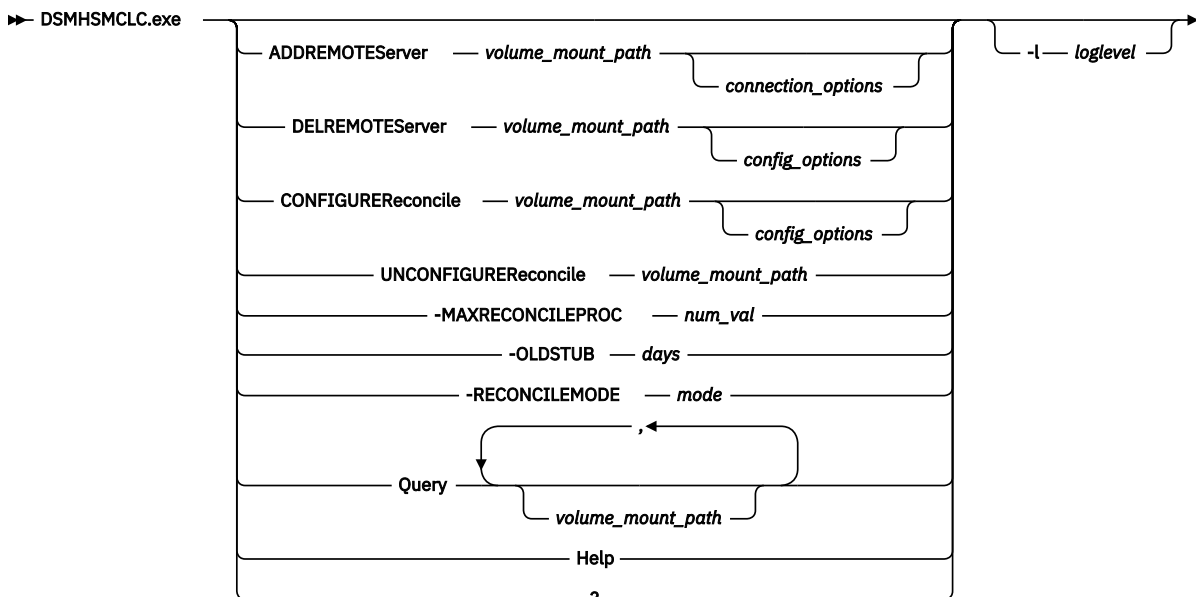
dsmhsmc1c.exe

Use the **dsmhsmc1c.exe** command to set and query the configuration of reconciliation and threshold migration. The settings will be used the next time a reconciliation process or a threshold migration process starts.

Managing reconciliation tasks

You can use the **dsmhsmc1c.exe** command to configure reconciliation. You can configure reconciliation, deactivate reconciliation, and set the limit for reconciliation processes. You can also choose whether or not reconciliation processing deletes obsolete objects.

Syntax



Options

ADDREMOTEserver

Use this option to add a remote IBM® Storage Protect server to a reconciliation task. Before you can add a remote server, you must configure the volume for reconciliation processing. If a file is recalled when in moving state, the migrated object is not automatically deleted on the remote IBM® Storage Protect server. The migrated object remains on the remote IBM® Storage Protect server until the remote IBM® Storage Protect server is added to a reconciliation process, and the reconciliation process is run.

DELETEREMOTEserver

Use this option to delete a remote IBM® Storage Protect server connection from a reconciliation task.

CONFIGUREREconcile

Use this option to configure reconciliation for the specified volume mount path.

UNCONFIGUREREconcile

Use this option to remove reconciliation from the specified volume mount path. When you specify this option, reconciliation is deactivated and all configuration values are erased.

-MAXRECONCILEPROC num_val

Use this option to configure the number of reconciliation tasks that can run at the same time. If this number is reached, any additional reconciliation tasks are delayed until the running reconciliation task finishes. Specify a value from 1 to 16. The default is 3.

-OLDSTUB days

Use this option to record the number of old stub files. The reconciliation process counts the number of stub files on the file system that are at least as old as the age that you specify. After the next reconciliation process, the hsmmonitor.log file contains the number of stub files that are at least as old as the age that you specify. A trace record from the log file looks like this example:

```
I: Number of old/unused stubs (age > 400 days): 13467
```

Specify an age in days. The default value is 0 days. If the value is 0, the number of old stub files is not recorded.

-RECONCILEMODE mode

Use this option to choose whether or not the reconciliation process deletes obsolete objects. If you do not specify *mode*, the command displays the current value of *mode*. After you change the value of *mode*, you must restart the HSM monitor service. If you specify *mode*, it must be one of the following values:

NORMAL

The reconciliation process marks, unmarks, and deletes objects on the IBM® Storage Protect server. The marked, unmarked, and deleted objects are recorded in a list file. The list file name is `hsmmonitor-delete-YYYYMMDD-hhmmss.log`, where `YYYYMMDD` indicates the date and `hhmmss` indicates the time when the HSM monitor service was started.

EMULATION

The reconciliation process runs in emulation mode. The reconciliation process does not mark, unmark, or delete objects on the IBM® Storage Protect server. The log output lists the objects that would be marked, unmarked, or deleted if the reconciliation process was run in normal mode. The objects are recorded in the `hsmmonitor-delete-YYYYMMDD-hhmmss.log` list file. `YYYYMMDD` indicates the date and `hhmmss` indicates the time when the HSM monitor service was started.

Query

Use this option to query the threshold migration configuration and reconciliation configuration of one or more volumes. Separate volume mount paths with a comma and no blank space. The default is all configured volumes.

In addition to configuration values, the query can display the following information for each volume, depending on whether threshold migration, reconciliation, or both, are configured:

- Time of next reconcile process
- Space usage
- Running processes:
 - Reconcile
 - Threshold migration (TM)
 - Scan for TM candidates
 - Validation of TM candidates

volume_mount_path

Specify the volume mount path. Because it is possible for a single volume to be mounted by more than one path, always specify a volume by the same mount path. Reconciliation, threshold migration, and migration jobs must all reference the volume by the same path.

connection_options

If the operation involves a remote file server, you must specify an IBM® Storage Protect connection.

You can specify a connection by specifying the two parts of a connection pair or by specifying a shortcut.

Specify the two parts of a connection

Use the **server** and **user** parameters:

-Server *TSM_host_name*

Specify the IBM® Storage Protect server part of a connection pair. The value of *TSM_host_name* is not case-sensitive. Specify *TSM_host_name* with the value of the **TCPSERVERADDRESS** option and the value of the **TCPPORT** option, separated with a colon. For example: `127.0.0.1:1500`

-User *node_name*

Specify the IBM® Storage Protect node part of a connection pair. Use the same value that you used to define the IBM® Storage Protect server connection. If the connection to the IBM® Storage Protect server was configured with the **asnodename** option, specify the value of the **asnodename** option. If the connection was configured without the **asnodename** option, specify the value of the **nodename** option. The value of *node_name* is not case-sensitive.

Specify a connection shortcut

Use the **connection** parameter:

-Connection *shortcut*

The *shortcut* value is one or two characters and is generated by the HSM for Windows™ client. Connection shortcuts include these examples:

- *l* (local)
- *r1* (remote connection 1)
- *r2* (remote connection 2)

config_options

You can specify any of the following configuration options. Specify each option no more than once. If the volume is not configured, omitting the option from the command configures the volume with the default

value for the option. If the volume is configured, omitting the option from the command leaves that configuration value unchanged.

-NEXTREConcile YYYY-MM-DD-hh-mm

Use this option to configure when the next regular reconciliation occurs. The date and time indicate year (YYYY), month (MM), day (DD), hour (hh), and minute (mm). Separate each element with a dash (-). The default is the current date and time.

-RECONCILEINTerval hours

Use this option to configure the number of hours between reconciliations. The interval starts when a reconciliation ends. If this option is set to 0, automatic reconciliation is deactivated. The range of acceptable values is 0-876000. The default is 720 hours.

-RECONCILENOW no|yes

Use this option to start reconciliation immediately. The default is no.

-RECONCILEPROTECTED no|yes

Set this option to reconcile protected files. A protected file is a file that was migrated and the file or stub file was deleted from the file system by a migration job. The default is no.

-RECONCILEPROTAGE days

When you set `reconcileprotected yes`, specify the time period as a number of days. The reconciliation process processes protected files that became protected only before the time period. The default is 1095 days.

-FILESPlaceList ALL | file space,file space

Use this option to configure the file spaces that are used when this volume is reconciled. Separate file space names with a comma and with no blank spaces. If you specify no file space names, or specify ALL, all available file spaces are used for reconciliation.

You can improve the reconciliation performance by restricting the list to the file spaces that contain migrated files of the volume that you are configuring.

-L loglevel

Specify the type of information that is to be recorded in logs and trace files. You can specify one or more values with no commas or blank space separators. Severe and error messages are always recorded. The default combination is severe, error, warning, information, and library (SEWIL). The following values are valid:

- C (event)
- D (debug)
- E (error)
- F (flush)
- I (information)
- K (driver)
- L (library)
- S (severe)
- T (trace)
- U (user)
- W (warning)
- X (dump)

Help

Use this option to display help for the command. Entering the command with no options also displays help for the command.

?

Use this option to display help for the command. Entering the command with no options also displays help for the command.

Examples

Task

Volume e: \ is not yet configured for reconciliation. Configure volume e: \ for reconciliation. Accept the default values for all parameters.

Command: `dsmhsmclic configurer e:\`

Task

Configure the next reconciliation to start at midnight on the 1 December 2011 using file spaces `fileSpaceA` and `fileSpaceC`, with an interval of one year (8760 hours) until the next reconciliation.

Command: `dsmhsmclic configurer e:\ -nextrec 2011-12-01-00-00 -filesp fileSpaceA,fileSpaceC -reconcileint 8760`

Task

Volume `f:\` is not yet configured for reconciliation. Configure volume `f:\` for reconciliation. Accept the default values for all parameters except **reconcileinterval**.

Command: `dsmhsmclic configurer f:\ -reconcileinterval 1000`

Task

Volume `g:\` is already configured for reconciliation. Change only the **reconcileinterval** value for this volume.

Command: `dsmhsmclic configurer g:\ -reconcileint 800`

Task

Limit reconciliation among all volumes to one reconciliation process at a time.

Command: `dsmhsmclic -maxreconcileproc 1`

Task

Deactivate automatic reconciliation but do not erase reconciliation configuration of volume `e:\`.

Command: `dsmhsmclic configurer e:\ -reconcileint 0`

Task

Deactivate reconciliation and erase reconciliation configuration for volume `e:\`.

Command: `dsmhsmclic unconfigurer e:\`

Task

Add a remote server for reconciliation of volume `e:\`. The remote IBM® Storage Protect connection shortcut is `r1`.

Command: `dsmhsmclic addremotes e:\ -co r1`

Task

All stub files on volume `e:\` have been processed by the **hsmtasks** service and all obsolete objects have been removed from the remote IBM® Storage Protect server at `HAMBURG_TSM`. You want to remove the remote server from reconciliation of volume `e:\`. You cannot delete the connection file (`HAMBURG_TSM-TSMNODE.opt`) from the configuration directory because you still require the connection for reconciliation of other volumes.

Delete the remote server for reconciliation of volume `e:\`.

Command: `dsmhsmclic delremotes e:\ -se HAMBURG_TSM:1500 -us TSMNODE`

Task

Query the configuration of volumes `e:\` and `g:\`.

Command: `dsmhsmclic q e:\,g:\`

Task

Change the information that is recorded in log and trace files. Record dump and trace information, and (by default) severe and error information.

Command: `dsmhsmclic -l XT`

Task

Change the information that is recorded in log and trace files to the default.

Command: `dsmhsmclic -l`

Task

Display help for the **dsmhsmclic.exe** command (three methods are shown).

Command: `dsmhsmclic ?`

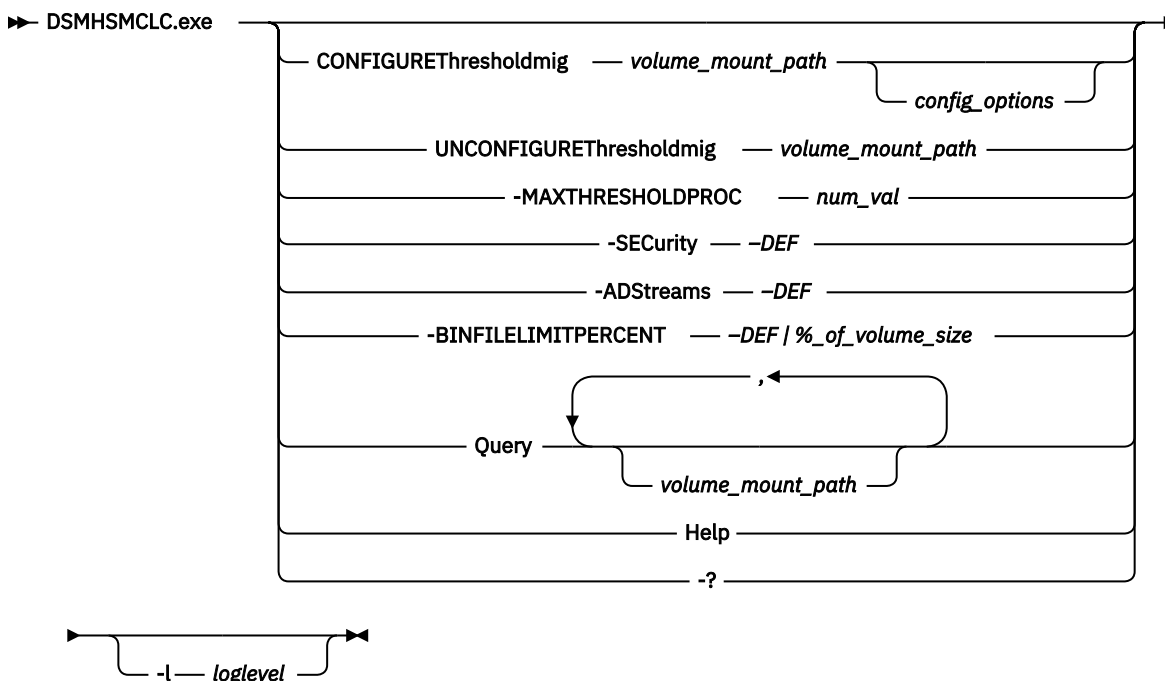
Command: `dsmhsmclic help`

Command: `dsmhsmclic`

Managing threshold migration

You can use the **dsmhsmclc.exe** command to manage threshold migration. You can configure threshold migration, deactivate threshold migration, set the limit for threshold migration processes, and query the configuration.

Syntax



Options

CONFIGUREThresholdmig

Use this option to configure threshold migration for the specified volume mount path.

UNCONFIGUREThresholdmig

Use this option to remove threshold migration from the specified volume mount path. When you specify this option, threshold migration is deactivated and all configuration values are erased.

-MAXTHRESHOLDPROC num_val

Use this option to configure the number of migration tasks that can occur simultaneously. The option applies to migration, scan, and validation tasks on all volumes. If this number is reached, any pending migration tasks are delayed until one of the running tasks finishes. The range of acceptable values is 1-16. The default is 3.

-SECurity yes|no

Use this option to configure whether file security attributes are migrated when the file is migrated. The default is the value that you set in the initial configuration wizard.

-ADStreams no|yes

Use this option to configure whether Windows™ alternate data streams data is migrated when the file is migrated. The default is the value that you set in the initial configuration wizard.

-BINFILELIMITPERCENT percent_of_volume_size off|on

Use this option to limit the size of the candidate list file `dsmtreeashmig.bin`. By default, the option is set to off, and the optimization is disabled. A percentage value of 1 to 99 sets the optimization to on and limits the size of the candidate list file. For more information, see [“Small file size for threshold migration might require a cleanup” on page 118](#).

Query

Use this option to query the threshold migration configuration and reconciliation configuration of one or more volumes. Separate volume mount paths with a comma and no blank space. The default is all configured volumes.

In addition to configuration values, the query can display the following information for each volume, depending on whether threshold migration, reconciliation, or both, are configured:

- Time of next reconcile process
- Space usage
- Running processes:
 - Reconcile
 - Threshold migration (TM)
 - Scan for TM candidates
 - Validation of TM candidates

volume_mount_path

Specify the volume mount path. Because it is possible for a single volume to be mounted by more than one path, always specify a volume by the same mount path. Reconciliation, threshold migration, and migration jobs must all reference the volume by the same path.

config_options

You can specify any of the following configuration options. Specify each option no more than once. If the volume is not configured, omitting the option from the command configures the volume with the default value for the option. If the volume is configured, omitting the option from the command leaves that configuration value unchanged.

-FILESSpace file space

Use this option to configure the file space that is used for threshold migration.

On the initial configuration, you must specify a file space. After the initial configuration, this parameter is optional. Until you specify a different file space, files that are migrated from this volume are stored in this file space.

-MGMTclass management class

Use this option to configure the management class that is used for threshold migration of this volume. Specify an existing management class with an archive copy group, or specify DEFAULT to use the default management class of the active policy set. If the retention period of the selected management class is finite, a warning is issued.

-HIGHthreshold percent

Use this option to configure the disk usage that triggers when to start threshold migration. After the disk capacity reaches this percent of capacity, threshold migration begins. The range of acceptable values is 1-100. The default is 90.

-LOWthreshold percent

Use this option to configure the disk usage that triggers when to stop threshold migration. After the disk usage reaches this percent of capacity, threshold migration stops. The low threshold must be less than the high threshold. The range of acceptable values is 0-99. The default is 80.

-MONitorinterval minutes

Use this option to configure how frequently the HSM monitor service checks space usage on the disk. The time is measured in minutes. If the monitor interval is set to 0, monitoring is deactivated. The range of acceptable values is 0-9999. The default is 5.

-SCANinterval hours

Use this option to configure how frequently the HSM monitor service starts the file system scan to find migration candidates. The time is measured from the end of the last scan to the beginning of the next scan. The time is measured in hours. The range of acceptable values is 1-9999. The default is 24.

If a scan yields better quality candidates (older and larger files) than the previous scan, the interval is automatically decreased by a small amount. If a scan yields poorer quality candidates (newer and smaller files) than the previous scan, the interval is automatically increased by a small amount.

-CHECKCANDidatesinterval minutes

Use this option to configure how frequently the HSM monitor service validates the candidates in the candidates pool. The time is measured from the end of the last validation to the beginning of the next validation. The time is measured in minutes. If the interval is set to zero, validation is deactivated. The range of acceptable values is 0-9999. The default is 180.

-MINMIGFILESIZE kilobytes

Use this option to configure minimum file size for a valid migration candidate. The size is measured in kilobytes (KB). The range of acceptable values is 4-2147483647 (2 TB). The default is 4.

-MINMIGFILEAGE days

Use this option to configure minimum file age for a valid migration candidate. The age is measured in days. The range of acceptable values is 0-99999. The default is 360.

-MINAGETYPE Access|Create|Modify

Use this option to configure which time stamp is used to calculate the age of a file. Changing this option can make many files in the current pool of migration candidates no longer valid. The choices correspond to the file system time stamps for file creation, file modification, and file access. The default is the file access time.

-AGEWeight percent

Use this option to configure the importance of file age (relative to file size) when determining migration candidates.

The age weight and size weight of a file are computed relative to the configured minimum age and minimum size. Hence, a file that is twice as old as the minimum age has an age weight of 2. If the file is the minimum size, it has a size weight of 1.

When the importance of age relative to size is considered, the file's weight is computed in this way:
computed weight = (AGEWeight * (age weight)) + ((1-AGEWeight) * (size weight)).

For example, when AGEWeight = 50, the file has the same weight $((.5 * (2)) + ((1 - .5) * (1)) = 1.5)$ as a file that is only as old as the minimum age, but twice as big as the minimum size $((.5 * (1)) + (.5 * (2)) = 1.5)$. The weight of both files is 1.5.

If the AGEWeight option is not 50%, but 75%, the first file has a computed weight of 1.75 $((.75 * (2)) + ((1 - .75) * (1)) = 1.75)$, while for the younger but larger file, the computed weight is 1.25 $((.75 * (1)) + ((1 - .75) * (2)) = 1.25)$.

Specify a value from 0 to 100. The default is 50.

-BACKUPBEforemigrate yes|no

Use this option to configure whether migration requires backup. The default is the value that you set in the initial configuration wizard. If you use the **-backupbeforemigrate** option but do not specify yes or no, a file is backed up before it is migrated. The default is yes.

-OPTFILE options_file

Use this option to specify the options file for backup before migration. If you specify **-OPTFILE=DEFAULT**, the backup-archive client chooses the options file. The backup-archive client chooses the options file even if the volume was previously configured to use another options file. The backup-archive client chooses the options file even if you specified another options file in the initial configuration wizard.

-THRESHOLDMIGNOW yes|no

Use this option to configure an immediate threshold migration. If disk usage is greater than the low threshold, files are migrated until the low threshold is reached. The default is no.

-SCANNOW yes|no

Use this option to configure an immediate scan of the volume to find migration candidates. The default is no.

-L loglevel

Specify the type of information that is to be recorded in logs and trace files. You can specify one or more values with no commas or blank space separators. Severe and error messages are always recorded. The default combination is severe, error, warning, information, and library (SEWIL). The following values are valid:

C (event)
D (debug)
E (error)
F (flush)
I (information)
K (driver)
L (library)
S (severe)
T (trace)
U (user)
W (warning)
X (dump)

Help

Use this option to display help for the command. Entering the command with no options also displays help for the command.

?

Use this option to display help for the command. Entering the command with no options also displays help for the command.

Examples

Task

Volume e: \ is not yet configured for threshold migration. Configure volume e: \ for threshold migration. Accept the default values for all parameters. (The file space name must be specified on the initial configuration).

Command: dsmhsmclc configuret e:\ -files computer10

Task

Volume e: \ was configured with default values. Raise the high and low thresholds for volume e: \. Monitor the volume more frequently.

Command: dsmhsmclc configuret e:\ -high 95 -low 90 -monitor 2

Task

Volume e: \ was configured with default values. Change the importance of size (relative to age) when picking migration candidates. Scan the volume for new candidates immediately.

Command: dsmhsmclc configuret e:\ -agew 25 -scannow yes

Task

Immediately begin a migration of volume e: \. Continue migrating files until the disk usage is 40% of capacity.

Command: dsmhsmclc configuret e:\ -low 40 -migratenow yes

Task

Limit threshold migration among all volumes to one threshold migration process at a time.

Command: dsmhsmclc -maxthresholdproc 1

Task

Deactivate threshold migration but do not erase threshold migration configuration of volume e: \.

Command: dsmhsmclc configuret e:\ -monitorinterval 0

Task

Deactivate threshold migration and erase threshold migration configuration for volume e: \.

Command: dsmhsmclc unconfiguret e:\

Task

Set a new management class MC2 for files that are migrated from volume f: \ by threshold migration.

Command: dsmhsmclc configuret f:\ -mgmt MC2

Task

Query the configuration of volumes e: \ and g: \.

Command: dsmhsmclc q e:\,g:\

Task

Change the information that is recorded in log and trace files. Record dump and trace information, and (by default) severe and error information.

Command: dsmhsmclc -l XT

Task

Change the information that is recorded in log and trace files to the default.

Command: `dsmhsmc1c -l`

Task

Display help for the **dsmhsmc1c.exe** command (three methods are shown).

Command: `dsmhsmc1c ?`

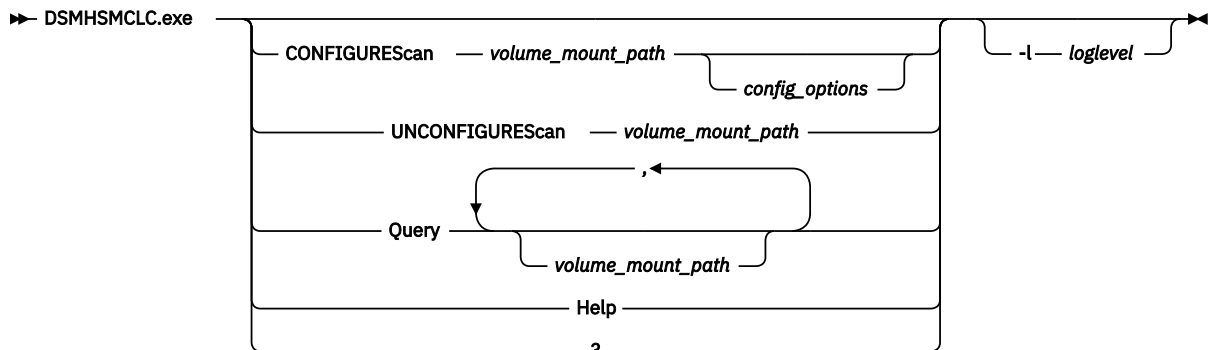
Command: `dsmhsmc1c help`

Command: `dsmhsmc1c`

Managing scans for volume statistics

You can use the **dsmhsmc1c.exe** command to manage scans for volume statistics. You can configure scans, set schedules for scans, deactivate scans, and query the configuration.

Syntax



Options

CONFIGUREScan

Use this option to configure a scan for volume statistics for the specified volume mount path.

UNCONFIGUREScan

Use this option to remove the scan for volume statistics from the specified volume mount path.

Query

Use this option to query the threshold migration configuration and reconciliation configuration of one or more volumes. Separate volume mount paths with a comma and no blank space. The default is all configured volumes.

In addition to configuration values, the query can display the following information for each volume, depending on whether threshold migration, reconciliation, scan for volume statistics, or all three, are configured:

- Time of next process
- Space usage
- Running processes:
 - Reconcile
 - Threshold migration (TM)
 - Scan for TM candidates
 - Validation of TM candidates
 - Scan for volume statistics

volume_mount_path

Specify the volume mount path. Because it is possible for a single volume to be mounted by more than one path, always specify a volume by the same mount path. Reconciliation, threshold migration, migration jobs, and volume scan jobs must all reference the volume by the same path.

config_options

You can specify any of the following configuration options. Specify each option no more than once. If the volume is not configured, omitting the option from the command configures the volume with the default value for the option. If the volume is configured, omitting the option from the command leaves that configuration value unchanged.

-NEXTScan YYYY-MM-DD-hh-mm

Use this option to configure when the next regular scan for volume statistics occurs. The date and time indicate year (YYYY), month (MM), day (DD), hour (hh), and minute (mm). Separate each element with a dash (-). The default is the current date and time.

-SCANinterval hours

Use this option to configure how frequently the HSM monitor service starts the scan for volume statistics. The time is measured in hours from the end of the last scan to the beginning of the next scan. You can specify a value in the range 1-9999. The default is 720.

-SCANNOW yes|no

Use this option to configure an immediate scan for volume statistics. The default is no.

-L loglevel

Specify the type of information that is to be recorded in logs and trace files. You can specify one or more values with no commas or blank space separators. Severe and error messages are always recorded. The default combination is severe, error, warning, information, and library (SEWIL). The following values are valid:

- C (event)
- D (debug)
- E (error)
- F (flush)
- I (information)
- K (driver)
- L (library)
- S (severe)
- T (trace)
- U (user)
- W (warning)
- X (dump)

Help

Use this option to display help for the command. Entering the command with no options also displays help for the command.

?

Use this option to display help for the command. Entering the command with no options also displays help for the command.

Example

Task

Configure volume T:\ to scan for volume statistics. Schedule a scan for once a week on Sundays at 6 am.

Command: dsmhsmclic configurescan t: -nextscan 2019-08-18-06-00 -scaninterval 168

The scan results are written to the scan log file hsmmonitor-volumes-stats.log.

Output

```
VNAME="T:\ " VSIZE="10734268416" VFREE="9105620992" VUSED="1628647424" VUSAGE="15.172412"
NFILES="240" NSTUBS="5"
```

```
FSIZEALL="1577975038" SSILOELOC="0" SSILOEFUL="8679424" SCANTYPE="scan"
SCANSTART="20190815-044636" SCANEND="20190815-044702"
```

- where:

VNAME="T:\ " volume name

VSIZE="10734268416" volume size (bytes)

VFREE="9105620992" amount of free space on the volume (bytes)


VUSED="1628647424" amount of used space on the volume (bytes)
 VUSAGE="15.172412" percentage of volume space that is used
 NFILES="240" number of files (stubs and resident)
 NSTUBS="5" number of stubs
 FSIZEALL="1577975038" total size of all files in the volume (bytes)
 SSIZELOC="0" total size of local stubs (bytes)
 SSIZEFUL="8679424" logical size of stubs (bytes)
 SCANTYPE="scan" scan type
 SCANSTART="20190815-044636" scan start date and time
 SCANEND="20190815-044702" scan end date and time

dsminfo.exe

Run the **dsminfo.exe** command from a **Command Prompt** window to view HSM for Windows™ client settings.

When you run this command the log file `dsminfo.log` is created.

Syntax



Options

info_options

You can specify any of the following options. Separate options with a space.

| Table 13: Options for dsminfo.exe | |
|-----------------------------------|--|
| Option | Description |
| all | Displays information for all options in this table |
| clclog | Displays the dsmdl.c.exe command log level |
| cluster | Displays cluster information |
| disk | Displays hard disks information |
| driver | Displays HSM for Windows® file-system driver version |
| errors | Displays only messages that contain installation errors |
| files | Displays all files of a valid HSM installation. |
| filter | Displays the attribute file filter and minimum file size |
| guilog | Displays the dsmgui.exe command log level |
| help | Displays the help for the options for this command |
| infolog | Displays dsminfo.exe command log level |
| installdir | Displays the installation directory |
| ip | Displays local computer IP addresses |
| mappings | Lists the hardware volume mappings |
| save | Saves the output to <code>check_installation.txt</code> (any further run of the command deletes this file) |
| servicelog | Displays the hsmervice.exe command log level |
| tivoli | Displays the versions of the IBM® Storage Protect backup-archive client and API |
| user | Displays the user name |
| version | Displays the HSM for Windows™ client version |
| win | Displays the Windows® version and fix pack |

| Option | Description |
|--------|--|
| wincp | Displays the Windows® default ANSI code page |

Help

Use this option to display help for the command. Entering the command with no options also displays help for the command.

Examples

Task

Display the version of the HSM for Windows™ client client.

Command: dsminfo version

Task

Display the logging level of the following commands: **hsmervice.exe**, **dsmgui.exe**, **dsmclc.exe**.

Command: dsminfo servicelog guilog clclog

Task

Display help for the **dsminfo.exe** command (two methods are shown).

Command: dsminfo help

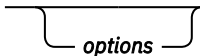
Command: dsminfo

dsmmove.exe

Run the **dsmmove.exe** command to move stub files to another location. If the other location is managed by a different IBM® Storage Protect server, the migrated file data is moved to the new IBM® Storage Protect server.

Run the **dsmmove.exe** command from a **Command Prompt** window on the local file server.

Syntax

► **DSMMOVE.exe**  *source_file_pattern* — *target_directory* ►

Parameters

options

You can specify any of the following options. Separate options with a blank space.

-d

Option **-d** specifies that stub files in retention state are moved. The retention is not restarted on the local IBM® Storage Protect server. By default, stub files in retention state are not moved. These files are considered deleted, but kept in the retention state by the IBM® Storage Protect server.

-f

Option **-f** specifies that a moved stub file replaces an existing file of the same name. You are not prompted to confirm replacement. By default, files are not replaced on the local file system and you are not prompted. A warning is logged.

-g *file_space*

Option **-g** specifies the file space in which the content of the stub files will be stored on the local IBM® Storage Protect server. You must specify option **-g** if files are moved to another file server. Do not specify this option if stub files are moved within a volume or to another volume on the same file server. For such moves, the migrated content of stub files on the local file server remains in the same file space.

-m *management_class*

Option **-m** specifies an IBM® Storage Protect server management class. If you do not specify this option, the moved stub files are bound to the default management class. Do not specify this option if stub files are moved within a volume or to another volume on the same file server. For such moves, no new object is created on the IBM® Storage Protect server.

-r

If option **-r** is specified, the **dsmove.exe** command traverses subdirectories on the remote file server when scanning for stub files to move. When reaching the volume boundary, the **dsmove.exe** command stops. The command does not traverse into nested volumes.

-s

If option **-s** is specified, the **dsmove.exe** command applies the security access control list (ACL) of the remote stub file to the local stub file.

This option does not affect the ACLs of local directory objects. The ACLs of the remote directory objects are not applied to the local directory objects.

connection_options

If the operation involves a remote file server, you must specify an IBM® Storage Protect connection.

You can specify a connection by specifying the two parts of a connection pair or by specifying a shortcut.

Specify the two parts of a connection

Use the **h** and **u** parameters:

-h TSM_host_name

Specify the IBM® Storage Protect server part of a connection pair. The value of *TSM_host_name* is not case-sensitive. Specify *TSM_host_name* with the value of the **TCPSERVERADDRESS** option and the value of the **TCPPORT** option, separated with a colon. For example: 127.0.0.1:1500

-u node_name

Specify the IBM® Storage Protect node part of a connection pair. Use the same value that you used to define the IBM® Storage Protect server connection. If the connection to the IBM® Storage Protect server was configured with the **asnodename** option, specify the value of the **asnodename** option. If the connection was configured without the **asnodename** option, specify the value of the **nodename** option. The value of *node_name* is not case-sensitive.

Specify a connection shortcut

Instead of specifying the host name and node name parts of a connection, you can specify a connection shortcut. Use the **c** parameter to specify a connection shortcut:

-c shortcut

The *shortcut* value is one or two characters and is generated by the HSM for Windows™ client. Connection shortcuts include these examples:

- *l* (local)
- *r1* (remote connection 1)
- *r2* (remote connection 2)

Tip: To display connection shortcuts, run the **dsmove** command with no parameters. Help for the command is displayed and defined shortcuts are displayed.

source_file_pattern

Specify the location of the stub files. You can use wildcard characters. If stub files are moved to another file server, you must use a UNC (Universal Naming Convention) path name.

You cannot move stub files with the **dsmove.exe** command if the host name of the remote file server is the same as the host name of the local file server. If the host name is the same, you can move a migrated file by recalling the file, moving the file, and then migrating the file again.

The **dsmove.exe** command does not traverse into nested volumes, even if option **-r** is specified. To move data from nested volumes, run the **dsmove.exe** command for each nested volume.

target_directory

Specify where to move the stub files. If the local directory does not exist, the stub moving tool creates the directory with default security settings.

Entering the command with no options displays help for the command. The help displays the command syntax and previously defined connection shortcuts.

Examples

Task

Move migrated files from remote directory `\\REMOTE_HOST\dir\` and all subdirectories to local directory `E:\new_dir`. Indicate the connection with the `-h` (*host_name*) and `-u` (*node_name*) parameters. Accept the default management class.

Command:

```
dsmove -h 123.456.789.1:1505 -u TSMNODE -g tmspace -r  
\\REMOTE_HOST\dir\* E:\new_dir
```

Task

Move migrated PDF files (*.pdf) from remote directory `\\REMOTE_HOST\proj1\` to local directory `F:\proj1\PDFs`. Indicate the connection with the `-c` (*shortcut*) parameter. The connection is assigned shortcut value `r2`.

Command:

```
dsmove -c r2 -g projects -m DEFAULT \\REMOTE_HOST\proj1\*.pdf  
F:\proj1\new_PDFs
```

Task

Move migrated files from local directory `G:\proj3\` and all subdirectories to local directory `F:\proj3\`. Directory `G:\proj3\` and directory `F:\proj3\` are on the same file server.

Command:

```
dsmove -r G:\proj3\* F:\proj1
```

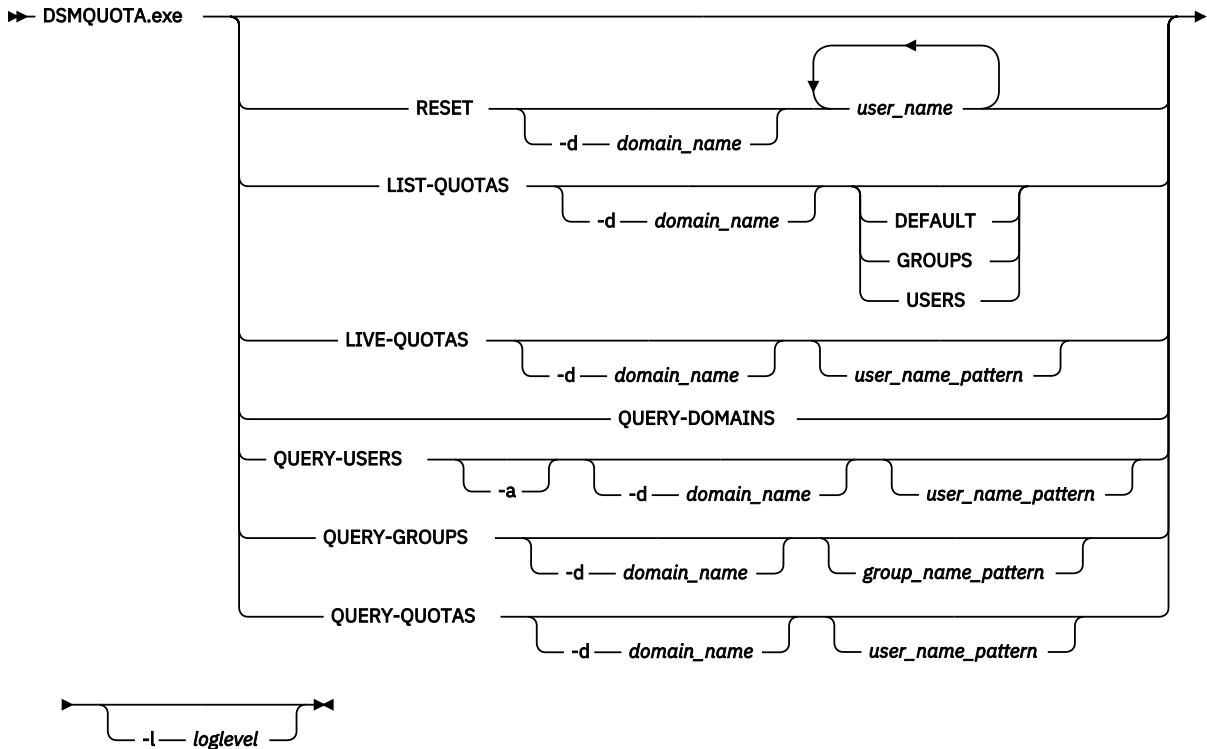
dsmquota.exe

Run the **dsmquota.exe** command to display user and group quotas or to reset the quota recall counter for one or more users.

Use the **dsmquota.exe** command to do the following quota tasks:

- Reset the quota recall counter for one or more user accounts
- List HSM quota definitions of user and group accounts
- Query Windows™ domains
- Query Windows™ user accounts
- Query Windows™ group accounts
- Evaluate effective quotas
- Display live quotas

Syntax



Parameters

RESET

Reset the quota recall counter for one or more user accounts. Separate user account names with a blank space. Each user account name is in the format *domain-name\user_name*. If you omit *domain-name*, the local host is used.

The domain (**-d**) parameter indicates that all user accounts are on the specified domain. After the domain parameter, list only the user account names. For example: `dsmquota reset -d domain1 user1 user2 user3`.

Restriction: If you use the domain parameter, you cannot reset quotas of user accounts on the local host.

LIST-QUOTAS

Display a list of quotas. You can restrict the list to a domain that you specify with the domain (**-d**) parameter. You can restrict the list to only user account quotas with the **users** parameter. You can restrict the list to only group account quotas with the **groups** parameter. You can restrict the list to only default quotas with the **default** parameter.

LIVE-QUOTAS

Display a list of user accounts that have a file recalls counter greater than 0. The record for each user account displays the number of file recalls in the period, and the file recalls quota in parentheses. You can restrict the list to a domain that you specify with the domain (**-d**) parameter.

Enter a value for *user_name_initial_chars* to filter group account names. The command displays all user accounts with a non-zero file recalls counter and that begin with the value.

QUERY-DOMAINS

Display a list of Windows™ domains.

QUERY-USERS

Display a list of Windows™ user accounts. You can restrict the list to a domain that you specify with the domain (**-d**) parameter.

Use the **-d** parameter to display more detailed information about user accounts, including group membership.

Enter a value for *user_name_initial_chars* to match user account names. The command displays all user account names that begin with the value.

QUERY-GROUPS

Display a list of Windows™ account groups. You can restrict the list to a domain that you specify with the domain (-d) parameter.

Enter a value for *group_name_initial_chars* to match group account names. The command displays all group account names that begin with the value.

QUERY-QUOTAS

Display the effective user account quota. The HSM for Windows™ client determines which user account quota, group account quota, and default quota definitions apply to a user account. The HSM for Windows™ client determines one effective quota for the user account.

The output displays quota information for each user account that matches the query:

- The user name
- The quota definition as file recalls per time span
- The effective quota type:
 - If the user account quota definition is the effective quota, the user account name is listed.
 - If the group account quota definition is the effective quota, the group account name is listed.
 - If the default quota definition is the effective quota, default quota is listed.

You can restrict the list to a domain that you specify with the domain (-d) parameter.

Enter *user_name_pattern* to match user account names. You can use wildcard character * to match one or more characters and ? to match one character. The command displays all user account names that match the pattern.

Example

Task

Reset the quota recall counter for local user accounts user43 and user78.

Command: `dsmquota reset user43 user78`

Task

Reset the quota recall counter for user accounts on different domains.

Command: `dsmquota reset domain5\user16 domain3\user56`

dsmtool.exe

Run the **dsmtool.exe** command to display the quantity, size, and expiration period of migrated objects in IBM® Storage Protect storage.

You can display the occupancy data of migrated files in the following ways:

- Use the IBM® Storage Protect administrative command **query occupancy**.
- The HSM for Windows™ client command **dsmclic listfilespace**s
- The HSM for Windows™ client command **dsmtool occupancy**

The administrative command **query occupancy** determines the amount of storage that a client is using on the IBM® Storage Protect server. The command displays the space that is used in an IBM® Storage Protect storage pool. The query can be refined by specifying a node name, file space name and type of data. You must have administrative access to the IBM® Storage Protect server to use this command.


The occupancy data that is displayed by the command **dsmclic listfilespace**s is the sum of file sizes of all migrated files for a file space. Occupancy also includes information for managing the migrated files.

Compression, data deduplication, and expirations on the IBM® Storage Protect server are not reflected in the statistics from the **dsmclic listfilespace**s command. The occupancy data is refreshed when you run the **dsmtool** command with the **occupancy** or **statistic** parameters.

For details on expirations due to copy group settings, see [Technote 1330160](#).

The **dsmtool occupancy** command displays the size that the migrated objects occupy on the file system. Compressed size and expired objects are not included in the occupancy calculation. Only uncompressed, unexpired objects sizes are calculated.

Syntax



Parameters

OCCUPANCY

Specify the **occupancy** option to display the number of migrated files and their total size. The size is calculated as the size of the resident file on the file system.

The size can be different from the administrative command **query occupancy**.

STATISTIC

Specify the **statistic** option to display the number of migrated files and their total size, and the versions and expiration periods of unexpired migrated files. The **statistic** option also displays the number and size of migrated files that are marked for deletion at the next reconciliation.

file_space_pattern

You can specify a file space. The specification can contain the wildcard character (*). If you do not specify a file space, the command displays information for all file spaces.

Example

Task

Display occupancy of all file spaces for the HSM for Windows™ client node.

Command: `dsmtool occupancy`

Task

Display the versions and expiration periods of migrated files on all file spaces that begin with *hsm*.

Command: `dsmtool statistic hsm*`

Troubleshooting the HSM for Windows™ client

You can diagnose and fix some common problems, such as those caused by antivirus software.

Troubleshooting steps and information

You can follow some general guidelines on troubleshooting and preparing information about IBM® Storage Protect HSM for Windows™ for IBM® support.

Trying the action again

1. Shut down the IBM Storage Protect™ HSM Recall Service.
2. Shut down the IBM Storage Protect™ HSM Tasks Service.
3. Shut down the IBM Storage Protect™ HSM Monitor Service, if it is installed.
4. Save and delete the log files.
5. Set the log levels to the highest level (**Full**) and ensure that the log file size is sufficiently large.
6. Restart the IBM Storage Protect™ HSM Recall Service (`hsmrecall.exe`) and verify that the service is running.
7. Restart the IBM Storage Protect™ HSM Tasks Service (`hsmtasks.exe`) and verify that the service is running.
8. Restart the IBM Storage Protect™ HSM Monitor Service (`hsmmonitor.exe`) and verify that the service is running.
9. Retry the action, if you still have an issue, retry the action using another method, for example:
 - Use the HSM for Windows™ client GUI instead of the **Command Prompt** window or vice versa.
 - Check permissions by creating a file in the directory of the stub file you are trying to retrieve.
 - From an application, such as MS Word, open and save the file in question.

Collecting data and files for IBM® support

A technical note provides steps for generating and collecting information that can help the IBM® Support Center assist you.

Offline stub files are recalled when they are first synchronized

Offline sub files are recalled the first time that Windows® synchronizes the offline files.

With the Windows® operating system, you can select a network file or folder to make it available offline. Windows® synchronizes your offline file with the network copy of the file when you reconnect to the network folder. The HSM for Windows™ client can migrate an offline file to IBM® Storage Protect storage. The first time that Windows® synchronizes the offline file, the HSM for Windows™ client recalls the migrated copy. The migrated copy is recalled even if you did not update a local copy after it was migrated to IBM® Storage Protect storage.

After the system synchronizes the copy, it does not recall the migrated copy the next time that synchronization takes place.

Problems with VSS during reconciliation

The HSM for Windows™ client uses VSS (Microsoft Volume Shadow Copy Service) during reconciliation. Errors can occur with VSS during reconciliation.

Look for clues to the VSS problem in the `msmmonitor-admin.log` file and in the `hsmmonitor.log` file.

For information about troubleshooting VSS problems with the IBM® Storage Protect backup-archive client, see [Troubleshooting: Using Windows Volume Shadow Copy Services](#).

Small migrated files occupy much space on IBM® Storage Protect server storage

Small files can occupy much space on IBM® Storage Protect server storage.

If you are using storage device class FILE on the IBM® Storage Protect server, the default minimum block size is 256 KB. Every file that is migrated occupies at least 256 KB in the storage pool. For example, with the default minimum block size, a 50 MB storage volume becomes full with 200 8-KB files.

You can eliminate the default minimum block size if you migrate to a storage pool that is defined with the attribute DATAFORMAT=NONBLOCK. You define storage pool attributes with the IBM® Storage Protect server command **DEFINE STGPOOL**.

Small file size for threshold migration might require a cleanup

If a small file size is configured for threshold migration, you might have to configure threshold migration to clean up of the candidate list file.

If the space savings that are reached during threshold migration are not as expected, check whether most of the migrated files on the volume are 4 KB - 6 KB in size. For those files, the space saving is negative if the files are listed in the candidate list file.

If this problem occurs, limit the size of the candidate list file `dsmthreshmig.bin` by using the **dsmhsmc1c** command option -BINFILELIMITPERCENT. When the specified file size is reached, the candidate list file is cleaned up. The cleanup includes deleting the existing candidate list file and creating a new, empty candidate list file. This way files already migrated are not included in the new candidate list file.

For more information, see [“Managing threshold migration” on page 104](#).

Accessibility features for the IBM® Storage Protect product family

Accessibility features assist users who have a disability, such as restricted mobility or limited vision, to use information technology content successfully.

Overview

The IBM® Storage Protect family of products includes the following major accessibility features:

- Keyboard-only operation
- Operations that use a screen reader

The IBM® Storage Protect family of products uses the latest W3C Standard, [WAI-ARIA 1.0 \(www.w3.org/TR/wai-aria/\)](http://www.w3.org/TR/wai-aria/), to ensure compliance with US Section 508 and [Web Content Accessibility Guidelines \(WCAG\) 2.0 \(www.w3.org/TR/WCAG20/\)](http://www.w3.org/TR/WCAG20/). To take advantage of accessibility features, use the latest release of your screen reader and the latest web browser that is supported by the product.

The product documentation in IBM® Documentation is enabled for accessibility.

Keyboard navigation

This product uses standard navigation keys.

Interface information

User interfaces do not have content that flashes 2 - 55 times per second.

Web user interfaces rely on cascading style sheets to render content properly and to provide a usable experience. The application provides an equivalent way for low-vision users to use system display settings, including high-contrast mode. You can control font size by using the device or web browser settings.

Web user interfaces include WAI-ARIA navigational landmarks that you can use to quickly navigate to functional areas in the application.

Vendor software

The IBM® Storage Protect product family includes certain vendor software that is not covered under the IBM license agreement. IBM makes no representation about the accessibility features of these products. Contact the vendor for accessibility information about its products.

Related accessibility information

In addition to standard IBM help desk and support websites, IBM has a TTY telephone service for use by deaf or hard of hearing customers to access sales and support services:

TTY service

800-IBM-3383 (800-426-3383)

(within North America)

For more information about the commitment that IBM has to accessibility, see [IBM Accessibility \(www.ibm.com/able\)](http://www.ibm.com/able).

Notices

This information was developed for products and services offered in the US. This material might be available from IBM® in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM® may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM® representative for information on the products and services currently available in your area. Any reference to an IBM® product, program, or service is not intended to state or imply that only that IBM® product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM® intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM® product, program, or service.

IBM® may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM® Director of Licensing

IBM® Corporation

North Castle Drive, MD-NC119

Armonk, NY 10504-1785

US

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM® Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing

Legal and Intellectual Property Law

IBM® Japan Ltd.

19-21, Nihonbashi-Hakozakicho, Chuo-ku

Tokyo 103-8510, Japan

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM® may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM® websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM® product and use of those websites is at your own risk.

IBM® may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM® Director of Licensing

IBM® Corporation

North Castle Drive, MD-NC119

Armonk, NY 10504-1785

US

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM® under terms of the IBM® Customer Agreement, IBM® International Program License Agreement or any equivalent agreement between us.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

Information concerning non-IBM® products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM® has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM® products. Questions on the capabilities of non-IBM® products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM®, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM®, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM® shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows: © (your company name) (year). Portions of this code are derived from IBM® Corp. Sample Programs. © Copyright IBM® Corp. _enter the year or years_.

Trademarks

IBM®, the IBM® logo, and ibm.com® are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM® or other companies. A current list of IBM® trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Adobe™ is a registered trademark of Adobe™ Systems Incorporated in the United States, and/or other countries.

Linear Tape-Open™, LTO™, and Ultrium™ are trademarks of HP, IBM® Corp. and Quantum in the U.S. and other countries.

Intel™ and Itanium™ are trademarks or registered trademarks of Intel™ Corporation or its subsidiaries in the United States and other countries.

The registered trademark Linux® is used pursuant to a sublicense from the Linux® Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

Microsoft™, Windows™, and Windows NT™ are trademarks of Microsoft™ Corporation in the United States, other countries, or both.

Java™ and all Java™-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Red Hat®, OpenShift®, Ansible®, and Ceph® are trademarks or registered trademarks of Red Hat®, Inc. or its subsidiaries in the United States and other countries.

UNIX® is a registered trademark of The Open Group in the United States and other countries.

VMware, VMware vCenter Server™, and VMware vSphere™ are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and/or other jurisdictions.

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM® website.

Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM®.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM®.

Rights

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM® reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM®, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM® MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Privacy policy considerations

IBM® Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM®'s Privacy Policy at <http://www.ibm.com/privacy> and IBM®'s Online Privacy Statement at <http://www.ibm.com/privacy/details> in the section entitled "Cookies, Web Beacons and Other Technologies," and the "IBM® Software Products and Software-as-a-Service Privacy Statement" at <http://www.ibm.com/software/info/product-privacy>.

Glossary

Index

36, 96, 97, 111, 110, 113, 115

backup of stubs 59

migration trigger 52

options

99, 99, 99, 99, 99, 99, 99, 99, 99, 99

(log level) 99

question mark () 99

overview 17

restoring default security attributes 61

when to restart 29, 29

(

(log level option)

command

parameter 77, 80, 80, 84, 78, 76, 82, 85, 89, 87, 90, 92, 94

reconciliation 99

threshold migration 108, 104

A

accessibility features 119

adding volumes

restart 29, 29

ADS

excluding by name 35

antivirus

stub file and recall 24

troubleshooting 24

B

Back up Migratable Files option 29

backing up and restoring migrated files 58

backing up migrated files

options 59

backup before migrate

choosing options file 58

backup options file

choosing 58

backup-archive client

backing up and restoring migrated files 58

controlling backups of migrated files 59

controlling location of temporary retrieves 59

controlling restores of migrated files 61

limiting temporary file copies 61

C

calculate migration savings 48

changing retention of migration copies 34

Check stub file reparse content option

overview 58

checkreparsecontent option

controlling backups of migrated files 59

closing idle server connections 41

cluster

installation 25

cluster environment

installation planning 23, 26

command

overview 76

parameter 77, 80, 80, 84, 78, 76, 82, 85, 89, 87, 90, 92, 94

reference 111, 113, 115

tasks 67

using 68

command line

return codes for operations 76

commands

96, 97, 111, 110, 113, 115

overview 76

parameter 77, 80, 80, 84, 78, 76, 82, 85, 89, 87, 90, 92, 94

reconciliation 99

tasks 67

threshold migration 108, 104

using 68

casing 75

minimum abbreviation 75

summary 75

using in executables 76

using in shell scripts 76

compatibility with other software 23

configuration

GUI 29

GUI

cluster environment 25, 25

connection to server 29

initial 29

reconciliation

command 99

GUI 65

threshold migration

command 108, 104

configuration file

setting location 36

configuration wizard 29

configuration wizard

cluster environment 25

configuring HSM client

to secondary server 32

configuring retention of migration copies 32

connection to server 29

connection to server

command 92

Connections parameter 35

ConnectionTimeout parameter 35

creating file spaces

command 76, 82

GUI 34

D

- date format setting [35](#)
- deleting migrated files from storage
 - command [78, 85](#)
- deleting obsolete quota entries [41](#)
- deleting protected files from storage [66](#)
- deletion results, displaying [71, 72](#)
- DirectoryAttributesFilter parameter [35](#)
- disability [119](#)
- display legends for table headers
 - command [80](#)
- display option defaults
 - command [77](#)
- displaying HSM results [71, 72](#)
- displaying listing files [71, 72](#)
- downgrade restrictions [24](#)
- drive letters
 - changing
 - reconciliation [64](#)
- dsmhsmclc.exe
 - options
 - [108, 108, 108, 108, 108, 108, 108, 108, 108, 108,](#)
[108, 108, 108, 108, 108, 108, 108, 108, 108,](#)
[104, 104, 104, 104, 104, 104, 104, 104, 104,](#)
[104, 104, 104, 104, 104, 104, 104, 104, 104](#)
 - (log level) [108, 104](#)

E

- effective quota [39](#)
- emulation mode
 - reconciliation [99](#)
- encryption
 - backup performance [60](#)
- exclude conditions
 - examples [45](#)
 - migration [43](#)
- excluding ADS names [35](#)
- executable file
 - return codes from [76](#)

F

- file group
 - creating [47](#)
 - editing [47](#)
 - overview [47](#)
- file locations [36](#)
- file name limitations [23](#)
- file quota
 - defining [37](#)
- file space
 - configure
 - GUI [34](#)

G

- GUI [29](#)

H

- hardware mapping

- overview [69](#)
- task [70](#)
- hsmmonitor.exe
 - when to restart [29](#)
- hsm.service.exe
 - when to restart [29](#)
- hsmtasks service [67](#)

I

- IBM Documentation [8](#)
- IBM Support Assistant [117](#)
- include conditions
 - examples [45](#)
 - migration jobs [43](#)
- installation
 - cluster environment
 - planning [23, 26](#)
 - network distribution [25](#)
- installation planning [23](#)

J

- jobs
 - migration
 - overview [43](#)
 - running [48](#)
 - removing unused stubs from a file system [49](#)

K

- keyboard [119](#)

L

- language setting [35](#)
- limitations
 - ADS [24](#)
 - file name [23](#)
- limiting temporary file copies [61](#)
- list file
 - settings
 - command [75](#)
 - GUI [41](#)
- list migration
 - overview [51](#)
 - running [87](#)
- listing files, displaying [71, 72](#)
- listing management class properties
 - command [84](#)
- listing migrated files
 - command [80](#)
- live quota [40](#)
- local server
 - definition [67](#)
- local file server
 - definition [67](#)
- local stub file
 - definition [67](#)
- log file
 - settings
 - command [75](#)
 - GUI [41](#)

- log level
 - command
 - parameter [77](#), [80](#), [80](#), [84](#), [78](#), [76](#), [82](#), [85](#), [89](#), [87](#), [90](#), [92](#), [94](#)
- configuring with GUI [36](#)
- option
 - reconciliation [99](#)
 - threshold migration [108](#), [104](#)
- window [36](#)

M

- management class**
 - configuring [34](#), [32](#)
- managing backups**
 - migrated files [61](#)
- managing temporary file copies** [61](#)
- manually retrieving files**
 - GUI [56](#)
- map hardware**
 - overview [69](#)
 - task [70](#)
- maximum connections parameter** [35](#)
- migrated files**
 - backup options [59](#)
- manually retrieving**
 - command [94](#)
 - GUI [56](#)
 - moving [67](#), [67](#), [68](#)
 - restore options [61](#)
- migrating a list of files**
 - command [87](#)
 - description [51](#)

- comparison with threshold migration [13](#)
- defining jobs [43](#)
- displaying results [71](#), [72](#)
- dsmclic.exe [48](#), [49](#)
- exclude conditions [43](#)
- include conditions [43](#)
- jobs [13](#)
- list [13](#)
- overview [13](#), [13](#)
- removing unused stubs from a file system [49](#)
- retention [17](#)
- run jobs from [48](#)
- running from command prompt [76](#)
- running jobs [48](#)
- scheduling job [49](#)
- space savings [48](#)
- threshold**
 - candidates [51](#)
 - command [108](#), [104](#)
 - comparison with migration jobs [13](#)
 - configuring with [108](#), [104](#)
 - migration triggers [52](#)
 - monitoring space usage [52](#)
- with other input [51](#)
- migration candidates**

reconciliation 64

- moving migrated files
- displaying results [71](#), [72](#)
- overview [67](#)

moving stub files [68](#)

MSCS cluster
installation planning 23, 26

msiexec 25

- national language environments [23](#)
- new features in version 8.1 [8](#)
- number format setting [35](#)

offline files [117](#)

option
108, 104
option, 108, 108, 108, 108, 108, 108, 108, 108, 108, 108,
108, 108, 108, 108, 108, 108, 108, 108, 99, 99, 99, 99,
99, 99, 99, 99, 99, 99, 104, 104, 104, 104, 104, 104,
104, 104, 104, 104, 104, 104, 104, 104, 104, 104, 104,
104

- options**
 - Back up Migratable Files [29](#)
 - checkreparsecontent**
 - controlling backups of migrated files [59](#)

- restoring files [61](#)
- restoremigstate** option
 - restoring files [61](#)
- restoring migrated files**
 - backup-archive client options [61](#)
- retention of migrated files [17](#)
- retrieving files**
 - command [94](#)
 - displaying results [71](#), [72](#)
 - GUI [56](#)
 - Windows alternate data stream (ADS) data [94](#)
- return codes for operations [76](#)
- running**
 - migration jobs [48](#)

S

- scan for migration candidates [51](#)
- search and retrieve files**
 - command [94](#)
 - GUI [56](#)
- security attributes**
 - restoring the default [61](#)
- selective recall [15](#)
- selective retrieve [15](#)
- server**
 - configuring connection**
 - cluster environment [25](#)
 - command [92](#)
 - GUI [29](#)
- shell scripts**
 - return codes from [76](#)
 - using commands in [76](#)
- Skip migrated files option**
 - overview [58](#)
- skitmigrated option**
 - controlling backups of migrated files [59](#)
- Staging Directory option**
 - overview [58](#)
- stagingdirectory option**
 - controlling location of temporary retrieves [59](#)
- Storage space**
 - small migrated files occupy much space [118](#)
- stub files**
 - backup [59](#)
 - move settings [37](#)
 - moving [67](#), [67](#), [68](#)
 - overview [15](#)
 - removing unused stubs from a file system [49](#)
- selective list recall**
 - command [90](#)
- selective recall**
 - command [89](#)

- Symantec Antivirus [24](#)
- system volume**
 - space management [55](#)
 - threshold migration [55](#)

T

- temporary files**
 - setting location [36](#)
- threshold migration**
 - candidates [51](#)
 - comparison with migration jobs [13](#)
 - configuration**
 - command [108](#), [104](#)
 - migration triggers [52](#)
 - monitoring space usage [52](#)
 - summary [51](#)
 - system volume [55](#)
- Threshold migration**
 - small file size might require a cleanup of the candidate list file [118](#)
- time zone setting [35](#)
- Timeout parameter [35](#)
- trace file**
 - settings**
 - command [75](#)
 - GUI [41](#)
- transparent recall [15](#)
- troubleshooting**
 - antivirus [24](#)
 - preview reconciliation deletions [66](#)
 - steps [117](#)

U

- Unicode setting [35](#)
- upgrade from HSM 7.1.1 and earlier [28](#)

V

- validation of migration candidates [51](#)
- view migration job results [48](#)
- volume**
 - changing drive letters**
 - reconciliation [64](#)
- VSS problems [117](#)

W

- weighting of migration candidates [51](#)
- window [65](#)
- Windows alternate data stream (ADS) data [28](#)
- Windows alternate data stream (ADS) data**
 - recall restriction [15](#)
 - retrieving [94](#)

© Copyright International Business Machines Corporation 2005, 2025

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

