

HPSS Management Guide

**High Performance Storage System,
version 10.3.0.0.0, 29 September 2023**

HPSS Management Guide

High Performance Storage System, version 10.3.0.0.0, 29 September 2023

Table of Contents

.....	xiii
1. HPSS configuration overview	1
1.1. Introduction	1
1.2. Starting the SSM GUI for the first time	1
1.3. HPSS configuration roadmap (new HPSS sites)	2
1.4. Initial HPSS startup roadmap (all sites)	3
1.5. Additional configuration roadmap (all sites)	4
1.6. Verification checklists (all sites)	4
1.6.1. Configuration checklists	4
1.6.2. Operational checklists	6
1.6.3. Performance checklist	6
2. Security and system access	8
2.1. Security services	8
2.1.1. Security services configuration	8
2.1.2. Security mechanisms	9
2.1.2.1. UNIX	9
2.1.2.2. UNIX and NIS or LDAP	9
2.1.2.3. Kerberos 5	10
2.1.2.4. LDAP	10
2.2. HPSS server security ACLs	12
2.3. SSM user security	13
2.4. Location policy	14
2.4.1. Configuring and updating a location policy	14
2.4.2. Location policy configuration window	14
2.4.3. Deleting a location policy	15
2.5. Restricting user access to HPSS	15
2.5.1. Overview	15
2.5.2. Configuration File	16
2.5.2.1. File Permissions	16
2.5.2.2. File Format	16
2.5.3. ra_verify_config	17
2.5.4. Restriction Types	17
2.5.5. Example Restrict Access File	17
3. Using SSM	18
3.1. The SSM System Manager	18
3.1.1. Starting the SSM System Manager	18
3.1.2. Tuning the System Manager RPC thread pool and request queue sizes	18
3.1.3. Labeling the System Manager RPC program number	19
3.2. Quick startup of hpssgui	20
3.3. Configuration and startup of hpssgui and hpssadm	21
3.3.1. Configuring the System Manager authentication for SSM clients	22
3.3.2. Creating the SSM user accounts	23
3.3.2.1. The hpssuser utility	23
3.3.2.2. SSM user authorization	23
3.3.2.3. User keytabs	24
3.3.3. SSM configuration file	26
3.3.3.1. login.conf	28

3.3.3.2. krb5.conf (for use with Kerberos authentication only)	28
3.3.4. SSM help files (optional)	29
3.3.5. SSM desktop client packaging	29
3.3.5.1. Automatic SSM client packaging and installation	29
3.3.5.2. Manual SSM client packaging and installation	30
3.3.6. Using SSM through a firewall	31
3.3.6.1. The firewall problem	31
3.3.6.2. Solutions for operating through a firewall	31
3.3.6.3. Example: Using hpssgui through a firewall	32
3.4. Multiple SSM sessions	34
3.5. SSM window conventions	34
3.6. Common window elements	36
3.7. Help menu overview	39
3.8. Monitor, Operations, and Configure menus overview	40
3.8.1. Monitor menu	40
3.8.2. Operations menu	42
3.8.3. Configure menu	43
3.9. SSM-specific windows	44
3.9.1. HPSS Login	45
3.9.2. About HPSS	46
3.9.3. HPSS Health and Status	47
3.9.3.1. SM server connection status indicator	48
3.9.3.2. HPSS status	48
3.9.3.3. HPSS statistics	50
3.9.3.4. Menu tree	50
3.9.3.5. File menu	51
3.9.3.6. View menu	51
3.9.4. SSM information windows	52
3.9.4.1. System Manager Statistics window	52
3.9.4.2. User Session Information window	57
3.10. SSM list preferences	58
4. Global and subsystem configuration	62
4.1. Global Configuration window	63
4.2. Storage subsystems	66
4.2.1. Subsystems list window	67
4.2.2. Creating a new storage subsystem	68
4.2.3. Storage Subsystem Configuration window	69
4.2.3.1. Create storage subsystem metadata	71
4.2.3.2. Create storage subsystem configuration	72
4.2.3.3. Create storage subsystem servers	72
4.2.3.4. Assign a Gatekeeper if required	73
4.2.3.5. Assign storage resources to the storage subsystem	73
4.2.3.6. Create storage subsystem fileset and junction	73
4.2.3.7. Migration and purge policy overrides	73
4.2.3.8. Storage class threshold overrides	74
4.2.4. Modifying a storage subsystem	74
4.2.5. Deleting a storage subsystem	74
4.3. HPSS trashcans	75
4.3.1. The trashcan feature is on	76

4.3.2. Trashcan operations for the administrator	78
5. HPSS servers	80
5.1. Server list	80
5.2. Server configuration	85
5.2.1. Common server configuration	88
5.2.1.1. Basic controls	88
5.2.1.2. Execution controls	90
5.2.1.3. Interface controls	92
5.2.1.4. Security controls	94
5.2.1.5. Audit policy	96
5.2.2. Core Server-specific configuration	97
5.2.2.1. Additional Core Server configuration	101
5.2.2.2. Core Server I/O limit configuration	103
5.2.2.3. Core Server max open files configuration	105
5.2.3. Gatekeeper-specific configuration	106
5.2.4. Location Server additional configuration	107
5.2.5. Migration/Purge Server (MPS)-specific configuration	108
5.2.6. Mover-specific configuration	110
5.2.6.1. Mover-specific configuration window	110
5.2.6.2. Additional Mover configuration	112
5.2.7. Physical Volume Repository (PVR)-specific configuration	117
5.2.7.1. Operator PVR-specific configuration window	118
5.2.7.2. AML PVR-specific configuration window	120
5.2.7.3. SCSI PVR-specific configuration window	125
5.2.7.4. STK PVR-specific configuration window	129
5.2.8. RAIT Engine-specific configuration	134
5.2.8.1. RAIT Engine-specific configuration window	134
5.2.9. Deleting a server configuration	136
5.3. Monitoring server information	137
5.3.1. Basic Server Information	137
5.3.2. Specific server information	140
5.3.2.1. Core Server Information window	141
5.3.2.2. Gatekeeper Information window	153
5.3.2.3. Location Server Information window	156
5.3.2.4. Migration/Purge Server Information window	157
5.3.2.5. Mover Information window	158
5.3.2.6. Physical Volume Library (PVL) information window	159
5.3.2.7. Operator PVR information window	160
5.3.2.8. AML PVR information window	162
5.3.2.9. SCSI PVR information window	165
5.3.2.10. STK PVR information window	168
5.3.2.11. RAIT Engine Information window	171
5.4. Real-Time Monitoring (RTM)	172
5.4.1. RTM Summary List	173
5.4.2. RTM Detail	174
5.5. I/O Abort	182
5.5.1. Aborting an I/O	182
5.5.2. Expectations in Aborting I/O	183
5.6. Starting HPSS	184

5.6.1. Starting HPSS prerequisite software	184
5.6.2. Starting HPSS servers	184
5.6.2.1. Starting the Startup Daemons	184
5.6.2.2. Starting SSM	184
5.6.2.3. Starting other HPSS servers	185
5.6.2.4. Automatic server restart	185
5.7. Stopping HPSS	186
5.7.1. Shutting down an HPSS server	186
5.7.2. Shutting down all HPSS servers	187
5.7.3. Halting an HPSS server	187
5.7.4. Shutting down the SSM server	187
5.7.5. Shutting down the Startup Daemon	188
5.7.6. Stopping the prerequisite software	188
5.8. Server repair and reinitialization	188
5.8.1. Repairing an HPSS server	188
5.8.2. Reinitializing a server	189
5.9. Forcing an SSM connection	191
6. Storage configuration	192
6.1. Storage classes	192
6.1.1. Configured Storage Classes window	192
6.1.2. Disk Storage Class Configuration	194
6.1.3. Tape Storage Class Configuration	199
6.1.4. Storage class subsystem thresholds	204
6.1.4.1. Disk Storage Subsystem-Specific Thresholds	205
6.1.4.2. Tape Storage Subsystem-Specific Thresholds	207
6.1.5. Changing a storage class definition	208
6.1.6. Deleting a storage class definition	210
6.2. Storage hierarchies	210
6.2.1. Hierarchies window	210
6.2.2. Storage Hierarchy Configuration window	212
6.2.3. Changing a storage hierarchy definition	213
6.2.4. Deleting a storage hierarchy definition	213
6.3. Classes of Service	214
6.3.1. Classes of Service window	214
6.3.2. Class of Service Configuration window	216
6.3.3. Changing a Class of Service definition	222
6.3.4. Deleting a Class of Service definition	223
6.3.5. Changing a file's Class of Service	223
6.3.6. Canceling a Class of Service change request	223
6.4. Migration policies	224
6.4.1. Migration Policies window	224
6.4.2. Migration policy configuration	226
6.4.2.1. Disk Migration Policy configuration	226
6.4.2.2. Tape Migration Policy configuration	230
6.4.2.3. Changing a migration policy	232
6.4.2.4. Deleting a migration policy	232
6.5. Purge policies	233
6.5.1. Purge Policies window	233
6.5.2. Purge Policy configuration	234

6.5.3. Changing a purge policy	236
6.5.4. Deleting a purge policy	237
6.5.5. Purge parallelism	237
6.6. File families	238
6.6.1. File Family Configuration	239
6.6.2. Changing a file family	239
6.6.3. Deleting a file family	239
7. Device and drive management	240
7.1. Configure a new device and drive	240
7.1.1. Devices and Drives window	250
7.1.2. Enable variable block sizes for tape devices	255
7.1.3. Changing a drive's configuration	255
7.1.4. Deleting a drive's configuration	256
7.2. Monitoring devices and drives	257
7.2.1. Mover Device Information window	258
7.2.2. PVL Drive Information window	264
7.3. Changing device and drive state	268
7.3.1. Unlocking a drive	269
7.3.2. Locking a drive	269
7.3.3. Repairing the state of a device or drive	269
7.3.4. Resetting drive statistics	270
7.3.5. Tape Drive Quota list window	270
8. Volume and storage management	272
8.1. Adding storage space	272
8.1.1. Importing volumes into HPSS	272
8.1.1.1. Import Tape Volumes window	275
8.1.1.2. Selecting import type for tape cartridges	279
8.1.1.3. Import Disk Volumes window	280
8.1.1.4. Selecting import type for disk volumes	283
8.1.2. Creating storage resources	284
8.1.2.1. Create Tape Resources window	285
8.1.2.2. Create Disk Resources window	289
8.2. Removing storage space	292
8.2.1. Deleting storage resources	292
8.2.1.1. Rules for deleting resources	293
8.2.1.2. Delete Resources window	293
8.2.2. Exporting volumes from HPSS	296
8.2.2.1. Rules for exporting volumes	296
8.2.2.2. Export Volumes window	297
8.3. Monitoring storage space	300
8.3.1. Active Storage Classes window	300
8.3.2. MPS Disk Storage Class Information	305
8.3.3. MPS Tape Storage Class Information	310
8.4. Dealing with a space shortage	313
8.4.1. Manually starting migration	313
8.4.2. Manually starting purge	314
8.4.3. Repacking and reclaiming volumes	314
8.4.4. Verifying data integrity	316
8.4.4.1. Repack Virtual Volumes window	317

8.4.4.2. Reclaim Virtual Volumes window	319
8.5. Volume management	320
8.5.1. Lookup Cartridges and Volumes window	320
8.5.2. PVL Volume Information window	321
8.5.3. PVR Cartridge Information window	323
8.5.4. Core Server Volume and Segment windows	328
8.5.4.1. Core Server Disk Volume information window	328
8.5.4.2. Core Server Tape Volume information window	333
8.5.5. Changing Core Server volume condition	341
8.5.6. Moving PVR cartridges to another PVR	342
8.5.6.1. Move Cartridges window	343
8.6. Monitoring and managing volume mounts	346
8.6.1. PVL Job Queue window	346
8.6.2. PVL Request Information window	350
8.6.3. Canceling queued PVL requests	353
8.6.4. Tape Check-In Requests window	353
8.6.5. Tape Mount Requests window	355
8.6.6. Administrative tape dismounts	356
8.7. Storage technology replacement	356
9. Logging and status	359
9.1. Logging overview	359
9.2. Log policies	360
9.2.1. Creating a log policy	360
9.2.2. Logging Policies window	361
9.2.2.1. Logging Policy configuration window	362
9.2.3. Changing a log policy	363
9.2.4. Deleting a log policy	363
9.3. Managing syslog	364
9.3.1. Configuring syslog	364
9.3.2. Recommendations for syslog configuration	364
9.3.3. Interactions between HPSS and syslog	365
9.3.4. Configure syslog message size	365
9.3.5. Configure multi-part messages	366
9.3.6. Restart Syslog After Configuration Change	366
9.3.7. Managing rate-limiting and duplication	367
9.3.8. Considerations For SELinux	367
9.3.9. Managing log archiving	367
9.3.10. Recommendations for logrotate configuration	369
9.3.11. Viewing HPSS logs	370
9.4. Managing SSM alarms and events	370
9.4.1. Alarms and Events window	371
9.4.2. Alarm/Event Information	372
9.4.3. Diagnosing HPSS problems with alarms and events	373
9.4.4. Controlling SSM log message handling	374
9.4.4.1. Controlling the System Manager log message cache	374
9.4.4.2. Controlling log messages displayed by hpssgui and hpssadm	375
9.4.5. Media Access Logging	376
9.5. Log files of various types - a comprehensive list	377
9.6. Disk Mover & VV I/O Metrics	381

9.6.1. Sample Log of Disk I/O Operations	383
10. Filesets and junctions	385
10.1. Filesets & Junctions List	385
10.2. Creating an HPSS fileset	387
10.2.1. Create Fileset window	388
10.3. Managing existing filesets	389
10.3.1. Core Server Fileset Information window	390
10.4. Deleting filesets	393
10.5. Creating a junction	393
10.5.1. Create Junction window	393
10.6. Deleting a junction	394
11. Files, directories and objects by SOID	395
11.1. Files and Directories window	395
11.1.1. File/Directory Information window	396
11.2. Objects by SOID window	402
12. Tape aggregation	404
12.1. Overview of tape aggregation	404
12.2. Tape aggregation performance considerations	404
12.3. Ordered Migration and Full Aggregate Recall	405
12.3.1. Ordered migration example - read-optimized aggregates	406
12.4. Configuring tape aggregation	406
12.5. Debugging tape aggregation	406
12.6. Caveats	407
13. Tape Ordered Recall (TOR)	408
13.1. Overview of TOR	408
13.2. TOR ordering behavior	408
13.3. TOR performance considerations	409
13.4. Configuring TOR	410
13.5. Caveats	411
14. User accounts and accounting	412
14.1. Managing HPSS users	412
14.1.1. Adding HPSS users	412
14.1.1.1. Add all user ID types	412
14.1.1.2. Add a user ID	413
14.1.1.3. Add an SSM user ID	413
14.1.2. Deleting HPSS users	414
14.1.3. Listing HPSS users	414
14.1.4. Create an SSM client package	415
14.2. Accounting	415
14.2.1. Accounting Policy window	416
14.2.2. Accounting reports and status	418
14.2.2.1. Generating an accounting report	418
14.2.2.2. Accounting Status window	419
14.2.2.3. Interpreting the accounting report	420
14.2.3. Accounting procedures	423
14.2.3.1. Site-defined accounting configuration files and procedures	423
14.2.3.2. Accounting intervals and charges	425
15. User interfaces	427
15.1. Client API configuration	427

15.2. FTP/PFTP daemon configuration	429
16. Backup and recovery	435
16.1. HPSS metadata backup and recovery	435
16.1.1. HPSS administrator responsibilities for DB2	435
16.1.2. Overview of the DB2 backup process	436
16.1.2.1. Configuring DB2 for online backup	437
16.1.3. Overview of the DB2 recovery process	438
16.2. HPSS system environmental backup	438
16.2.1. HPSS file system backup	438
16.2.2. Operating system backup	439
16.2.3. Kerberos backup	439
16.3. HPSS user data recovery: Recovering HPSS files from damaged HPSS volumes	440
16.4. DB2 monitoring	442
16.5. DB2 space shortage	442
16.5.1. DMS table spaces	442
16.5.2. SMS table spaces	443
17. Management tools	444
17.1. Utility overview	444
17.1.1. Fileset and junction management	444
17.1.2. Tape management	444
17.1.3. System info	445
17.1.4. System management	446
17.1.5. User interfaces	447
17.1.6. Testing and debugging	447
17.1.7. Unsupported tools	447
17.2. Tape queue monitoring	448
17.2.1. Example	449
17.2.2. JSON field descriptions	450
17.3. HPSS ACL permission exceptions	453
A. Glossary of terms and acronyms	456
B. References	466
C. Developer acknowledgments	468

List of Tables

3.1. SSM general options	26
3.2. HPSSGUI-specific options	27
3.3. HPSSADM-specific options	27
5.1. Mover TCP pathname options	111
5.2. Auto Restart Count values	186
5.3. Server reinitialization behavior	189
7.1. Recommended settings for tape devices	247
8.1. Tape import types	279
8.2. Disk import types	283
9.1. Log file formats	368
11.1. Attributes of files and directories	401
15.1. Banner keywords	432

List of Examples

6.1. Max Active File Delay values	231
---	-----

Copyright notification. Copyright © 1992-2023 International Business Machines Corporation, The Regents of the University of California, Los Alamos National Security, LLC, Lawrence Livermore National Security, LLC, Sandia Corporation, and UT-Battelle.

All rights reserved.

Portions of this work were produced by Lawrence Livermore National Security, LLC, Lawrence Livermore National Laboratory (LLNL) under Contract No. DE-AC52-07NA27344 with the U.S. Department of Energy (DOE); by the University of California, Lawrence Berkeley National Laboratory (LBNL) under Contract No. DE-AC02-05CH11231 with DOE; by Los Alamos National Security, LLC, Los Alamos National Laboratory (LANL) under Contract No. DE-AC52-06NA25396 with DOE; by Sandia Corporation, Sandia National Laboratories (SNL) under Contract No. DE-AC04-94AL85000 with DOE; and by UT-Battelle, Oak Ridge National Laboratory (ORNL) under Contract No. DE-AC05-00OR22725 with DOE. The U.S. Government has certain reserved rights under its prime contracts with the Laboratories.

DISCLAIMER. Portions of this software were sponsored by an agency of the United States Government. Neither the United States, DOE, The Regents of the University of California, Los Alamos National Security, LLC, Lawrence Livermore National Security, LLC, Sandia Corporation, UT-Battelle, nor any of their employees, makes any warranty, express or implied, or assumes any liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights.

Trademark usage. High Performance Storage System is a trademark of International Business Machines Corporation.

IBM is a registered trademark of International Business Machines Corporation.

IBM, DB2, DB2 Universal Database, AIX, RISC/6000, pSeries, and xSeries are trademarks or registered trademarks of International Business Machines Corporation.

UNIX is a registered trademark of the Open Group.

Linux is a registered trademark of Linus Torvalds in the United States and other countries.

Kerberos is a trademark of the Massachusetts Institute of Technology.

Java is a registered trademark of Oracle and/or its affiliates.

ACSLs is a trademark of Oracle and/or its affiliates.

Microsoft Windows is a registered trademark of Microsoft Corporation.

Other brands and product names appearing herein may be trademarks or registered trademarks of third parties.

Who should read this book. The *HPSS Management Guide* is intended as a resource for HPSS administrators. For those performing the initial configuration for a new HPSS system, Chapter 1 provides a configuration roadmap. For both new systems and those upgraded from a previous release, Chapter 1 provides a configuration, operational, and performance checklist which should be consulted before bringing the system into production. The remaining chapters contain the details for configuring, reconfiguring, monitoring, and managing an HPSS system.

Conventions used in this book. Example commands that should be typed at a command line will be preceded by a percent sign ("%") and be presented in a Courier font:

```
% sample command
```

Any text preceded by a pound sign("#") should be considered comment lines:

```
# This is a comment
```

Angle brackets("<>") denote a required argument for a command:

```
% sample command <argument>
```

Square brackets("[]") denote an optional argument for a command:

```
% sample command [optional argument]
```

Vertical bars("|") denote different choices within an argument:

```
% sample command <argument1 | argument2>
```

A byte is an eight-bit data octet. A kilobyte, KB, is 1024 (2^{10}) bytes; a megabyte, MB, is 1,048,576 (2^{20}) bytes; a gigabyte, GB, is 1,073,741,824 (2^{30}) bytes; a terabyte, TB, is 1,099,511,627,776 (2^{40}) bytes; and a petabyte, PB, is 1,125,899,906,842,624 (2^{50}) bytes. An exabyte, EB, is 1,152,921,504,606,846,976 (2^{60}) bytes.

Chapter 1. HPSS configuration overview

1.1. Introduction

This chapter defines the high-level steps necessary to configure, start, and verify correct operation of a new 7.5 HPSS system, whether that system is created from scratch or created by conversion from any previous version of HPSS.

To create or modify the HPSS configuration, we recommend that the administrator first be familiar with the information described in the *HPSS basics* and *HPSS planning* chapters of the *HPSS Installation Guide*.

Before performing the procedures described in this chapter, be certain that the appropriate system preparation steps have been performed. See the *System preparation* chapter of the *HPSS Installation Guide* for more information. For a system created from scratch, be certain that the HPSS installation and infrastructure configurations have been completed. See the *HPSS installation and infrastructure configuration* chapter of the *HPSS Installation Guide* for more information. To convert from a previous system, see the *HPSS Conversion Guide* for the steps to follow.

1.2. Starting the SSM GUI for the first time

The HPSS system is ready to be configured using SSM once the HPSS software is installed on the node and the HPSS infrastructure components are configured. In order to start the SSM GUI you must first start all infrastructure components and the SSM System Manager as follows:

```
% /opt/hpss/bin/rc.hpss -m start
```

Next you will need to add an SSM Admin user. To do this you will need to invoke the **hpssuser** utility as follows:

```
% /opt/hpss/bin/hpssuser -add hpss -<unix|krb> -ssm
```



The above commands must be done as root.

Once the SSM Admin user has been created, you can invoke the SSM GUI as follows (for `hpssgui.pl` options, see the **hpssgui** man page):

```
% /opt/hpss/bin/hpssgui.pl
```

Note: *This command may be done as an HPSS user.*

When the SSM GUI is running you can begin to configure the rest of HPSS (servers, devices, and so on) as described in the following sections. For more information on SSM, refer to *Chapter 3, Using SSM*.

1.3. HPSS configuration roadmap (new HPSS sites)

The following steps summarize the configuration of an HPSS system from scratch (not upgrading from a previous release). It is important that the steps be performed in the order listed. Each step is required unless otherwise indicated. Each step is discussed in more detail in the referenced section.

1. Configure storage subsystems (see *Section 4.2.2, “Creating a new storage subsystem”*)



*Subsystems can be configured only partially at this time. The **Gatekeeper**, **Default COS**, and **Allowed COS** fields will be updated in a later step.*

2. Configure HPSS storage policies

- Accounting Policy (refer to *Section 14.2.1, “Accounting Policy window”*)
- Log Policies (refer to *Section 9.2, “Log policies”*)
- Location Policy (refer to *Section 2.4, “Location policy”*)
- Migration Policies (refer to *Section 6.4, “Migration policies”*)
- Purge Policies (refer to *Section 6.5, “Purge policies”*)

3. Configure HPSS storage characteristics

- Storage Classes (refer to *Section 6.1.1, “Configured Storage Classes window”*)
- Storage Hierarchies (refer to *Section 6.2, “Storage hierarchies”*)
- Classes of Service (refer to *Section 6.3, “Classes of Service”*)

4. Configure HPSS servers (refer to *Section 5.2, “Server configuration”*)

5. Create global configuration (refer to *Section 4.1, “Global Configuration window”*)

6. Configure MVR devices and PVL drives (refer to *Section 7.1, “Configure a new device and drive”*)

7. Configure file families, if used (refer to *Section 6.6, “File families”*)

8. Update storage subsystem configurations with Gatekeeper and COS information (refer to *Section 4.2.4, “Modifying a storage subsystem”* and *Section 4.2.3, “Storage Subsystem Configuration window”*)

9. Create the endpoint map (refer to *Section 5.2.4, “Location Server additional configuration”*).

1.4. Initial HPSS startup roadmap (all sites)

This section provides instructions for starting the HPSS servers and performing post-startup configuration. For sites which are converting from a previous version of HPSS, only Step 1 may be necessary. For sites configuring a new system from scratch, all steps are necessary:

1. Start the HPSS servers (refer to *Section 5.6.2, "Starting HPSS servers"*)
2. Unlock the PVL drives (refer to *Section 7.3.2, "Locking a drive"*)
3. Create HPSS storage space:
 - a. Import volumes into HPSS (refer to *Section 8.1.1, "Importing volumes into HPSS"*)
 - b. Create storage resources (refer to *Section 8.1.2, "Creating storage resources"*)
4. Create additional HPSS users (refer to *Section 14.1.1, "Adding HPSS users"*)
5. Create Filesets and Junctions (refer to *Section 10.1, "Filesets & Junctions List"* and *Section 10.5, "Creating a junction"*)
6. Create HPSS `/log` directory

If log archiving is enabled, using an HPSS name space tool such as **scrub** or **ftp**, create the `/log` directory in HPSS. This directory must be owned by `hpsslog` and have permissions `rwxr-xr-x`.

The `/log` directory can be created by the root user using **ftp** as follows:

```
% ftp <node> <HPSS Port> # login as root user
ftp> mkdir /log
ftp> quote site chown hpsslog /log
ftp> quote site chmod 755 /log
```

If Account Validation is enabled, principal "hpsslog" must appear in the ACCTVAL table with a default account code allowing write access to the `/log` directory. The utility **hpss_avaledit** can be used to verify and, if necessary, correct this configuration.

```
% hpss_avaledit -list
ACCOUNT ACCOUNT_NAME DEFAULT USER
...
301 hpsslog yes hpsslog
...
```

If "hpsslog" does not appear in the list, or if the DEFAULT value is "no", or the user name is something other than "hpsslog", do the following:

```
% hpss_avaledit -remove hpsslog -all # if necessary
% hpss_avaledit -add -uname hpsslog -aname hpsslog -def
```

1.5. Additional configuration roadmap (all sites)

This section provides a high level roadmap for additional HPSS configuration.

1. Configure HPSS user interfaces (refer to *Chapter 15, User interfaces*).
2. Set up backup for DB2 and other infrastructure (refer to *Chapter 16, Backup and recovery*).
3. Set up High Availability, if desired (see the *HPSS planning* chapter of the *HPSS Installation Guide*).
4. Optionally configure support for both authentication mechanisms (refer to the *Supporting both UNIX and Kerberos authentication for SSM* section of the *HPSS Installation Guide*, in the *HPSS installation and infrastructure configuration* chapter).

1.6. Verification checklists (all sites)

This section provides a number of checklists regarding configuration, operational, and performance issues.

1.6.1. Configuration checklists

After HPSS is running, the administrator should use the following checklists to verify that HPSS was configured or converted correctly.

Global configuration

- Verify that a **Default Class of Service** has been selected.
- Verify that a **Root Core Server** has been selected.

Storage subsystem configuration

- Verify that a **Default Class of Service** has been selected if desired.
- Verify that a **Gatekeeper** has been selected if gatekeeping or account validation is required.
- Verify that the **COS Name** list has been filled in correctly.
- Verify that a **Core Server** and **Migration/Purge Server** have been configured for each storage subsystem.
- Verify that each storage subsystem is accessible by using **lsjunctions** and ensuring that there is at least one junction to the Root fileset of each subsystem. (The root fileset for a given subsystem can be found in the specific configuration for the subsystem's Core Server)

Servers

- Verify that all required HPSS servers are configured and running.

- Verify that all servers are configured with the intended level of logging, whether using their server-specific policies or the default policy. Also verify that all Core Server and Mover log policies have the **DEBUG** flag turned on to aid in the diagnostics of any future problems.

Devices and drives

- Verify that all devices/drives are configured and each is assigned to an appropriate PVR or Mover.
- For tape devices, verify that the **Locate Support** option is enabled if supported.
- For tape devices, verify that the **NO-DELAY** option is enabled if supported by the device.
- For disk devices, verify that the **Capacity** and **Starting Offset** values are correct.
- Verify that all configured drives are unlocked.

Storage classes

- Verify that all storage classes are defined and each has sufficient free storage space.
- Verify that each storage class that will be migrated and purged is configured with the appropriate migration and purge policy.
- Verify that no storage class at the lowest level in a hierarchy is configured with a migration or purge policy.
- To support repack and recover of tape volumes, verify that the stripe width of each tape storage class is less than half of the number of available drives of the appropriate drive type.

Storage hierarchies

- Verify that all storage hierarchies are defined.

Classes of Service

- Verify that all Classes of Service are defined.
- Verify that each Class of Service (COS) is associated with the appropriate storage hierarchy.
- Verify that the COS is configured to use the characteristics of the hierarchy and the underlying storage classes. In addition, verify that the Classes of Service have the correct **Minimum File Size** and **Maximum File Size** values. If these sizes overlap, the file placement may be indeterminate when the user creates a file using the size hints. For Classes of Services which are not to be used as part of standard file placement, set their **Force Selection** flag to ON so that they will only be chosen if specified by their **COS ID**.
- Verify that Classes of Service with multiple copies have the **Retry Stage Failures from Secondary Copy** flag enabled.

File families, filesets, and junctions

- Verify that file families and filesets are created according to the site's requirements.

- Verify that each fileset is associated with the appropriate file family, COS, or both.
- Verify that each fileset has an associated junction.

User interfaces

- Verify that the desired HPSS user interfaces (such as FTP, PFTP, Client API, or any others) are configured (refer to *Chapter 15, User interfaces*).

1.6.2. Operational checklists

The administrator should follow these checklists to ensure that HPSS is operating properly.

Configured user interfaces

- Create files of various sizes on each defined COS.
- Verify that the files are created in the expected storage class with acceptable transfer rates.
- If necessary, redefine the associated storage class definition to enhance the throughput performance.
- Update the characteristics fields (such as **Transfer Rate**, **Latency**, and others) in the storage class and Class of Service definition to reflect actual performance results.
- After the files are created on the correct storage class, verify that the files are created with correct file ownerships and permissions.
- Verify that other file operations (such as **delete**, **chmod**, and **copy**) work.
- If accounting is configured, verify that the files are created with the correct accounting indices.
- If file families, filesets, and junctions are configured, verify that they work as intended.

Devices and drives

- Ensure that each drive mounts, reads, and writes.

Storage management

- Verify that migration and purge operations work as intended.
- Start Migration and Purge operations manually to verify that files are migrated and purged correctly.
- Verify that files can be accessed after being migrated/purged.
- Monitor free space from the top level storage class in each hierarchy to verify that the migration and purge policy are maintaining adequate free space.

1.6.3. Performance checklist

Measure data transfer rates in each COS for:

- Client writes to disk
- Migration from disk to tape
- Staging from tape to disk
- Client reads from disk

Transfer rates should be close to the speed of the underlying hardware. The actual hardware speeds can be obtained from their specifications and by testing directly from the operating system (for example, by using the Linux **dd** command to read and write to each device). Keep in mind that transfer performance can be limited by factors external to HPSS. For example, HPSS file read performance may be limited by the performance of the UNIX file system writing the file rather than limits inside HPSS.

Chapter 2. Security and system access

2.1. Security services

The approach to security within HPSS divides services into two APIs, known as mechanisms, each of which has multiple implementations. Configuration files control which implementation of each mechanism is used in the security realm (analogous to a DCE cell) for an HPSS system. Security mechanisms are implemented in shared object libraries and are described to HPSS by a configuration file. HPSS programs that need to use the mechanism dynamically link the library to the program when the program starts.

The first type of mechanism is the **authentication mechanism**. This API is used to acquire credentials and to verify the credentials of clients. Authentication verifies that a client really is who it claims to be.

The second type of mechanism is the **authorization mechanism**. Once a client's identity has been verified, this API is used to obtain the UID, GID, group membership, and other authorization details associated with the client that are used to determine the privileges accorded to the client and the resources to which it has access.

2.1.1. Security services configuration

Ordinarily, the configuration files that control HPSS's access to security services are set up either by the installation tool **mkhpss** or by the metadata conversion tools. This section is provided purely for reference.

- `/var/hpss`

When installing HPSS on a new system, **mkhpss** will create the **hpss** user with the home directory as `/var/hpss`. Beginning in 9.2, it becomes critical that the home directory point to the location of the HPSS configuration files, `/var/hpss` by default. If either the configuration directory is moved or the HPSS user has a different home directory, HPSS will fail to start. If possible, ensure that the *hpss* user home directory matches the location of the HPSS config data.

There is also an alternate method to specify this configuration home, for those who may be on a global user directory such as NIS or ActiveDirectory, and whose user homes cannot be customized. A symbolic link can be created from `/etc/hpss` to the location of the HPSS configuration directory base. For example:

```
ln -s /var/hpss /etc/hpss
```

Each of the files below is stored by default in `/var/hpss/etc`.

- `authz.conf`

This file defines which shared libraries provide implementations of the authorization mechanisms. It is a plain text file. Each line is either a comment beginning with `#` or consists of two fields

separated by whitespace: the path to a shared library and the name of the function used to initialize the security interface.

- `site.conf`

This file defines security realm options. This is a plain text file in which each line is a comment beginning with `#` or defines a combination of authentication and authorization for a realm. See the man page for `site.conf` for details.

2.1.2. Security mechanisms

HPSS 7.5 supports UNIX and Kerberos mechanisms for authentication. It supports LDAP and UNIX mechanisms for authorization.

2.1.2.1. UNIX

UNIX-based mechanisms are provided both for authentication and authorization. These can draw either from the actual UNIX user and group information on the current host or from a separately maintained set of files used only by HPSS. This behavior is controlled by the setting of the variable `HPSS_UNIX_USE_SYSTEM_COMMANDS` in `/var/hpss/etc/env.conf`. If this variable is set to any non-empty value other than `FALSE`, the actual UNIX user and group data will be used. Otherwise, local files created and maintained by the following HPSS utilities will be used. Consult the man pages for each utility for details of its use.

- **`hpss_unix_keytab`** is used to define "keytab" files that can be used to acquire credentials recognized by the UNIX authentication mechanism.
- **`hpss_unix_user`** is used to manage users in the HPSS password file (`/var/hpss/etc/passwd`).
- **`hpss_unix_group`** is used to manage users in the HPSS groups file (`/var/hpss/etc/group`).
- **`hpss_unix_passwd`** is used to change passwords of users in the HPSS password file.
- **`hpss_unix_keygen`** is used to create a key file containing a hexadecimal key. The key is used during UNIX authentication to encrypt keytab passwords. The encryption provides an extra layer of protection against forged passwords.

Keep in mind that the user and group databases must be kept synchronized across all nodes in an HPSS system. If using the actual UNIX information, this can be accomplished using a service such as NIS. If using the HPSS local files, these must manually be kept in synchronization across HPSS nodes.

2.1.2.2. UNIX and NIS or LDAP

When setting up HPSS on a system with network-based accounting systems such as NIS (Network Information Service) or LDAP (Lightweight Directory Access Protocol), it is still necessary to keep the HPSS server accounts local to the system. This is because the current keytab system for UNIX uses a hashing method to convert the account's password into a keytab. To test whether a keytab works, HPSS must be able to get the account's password hash from the `passwd` database and hash it to compare with the keytab hash. LDAP typically depends on authenticating to the LDAP server instead of conducting the comparison; hence, it does not even provide the hashes.

In addition to authentication difficulties, this also means that there are two locations for account information to be obtained. HPSS at present does not have a standard way to search "HPSS" local password files for some accounts while using LDAP for users. Because of this, the recommended approach for integrating NIS or LDAP into HPSS is the following:

- use UNIX authentication, and configured to use system files
- store HPSS server accounts and any users that will use keytabs (that is, with **hpssadm**) in the local system
- if using the HPSS accounts across multiple systems, copy the HPSS password and shadow entries and keytabs manually from the core system to its ancillary systems in order to maintain the same hashes.

2.1.2.3. Kerberos 5



The capability to use MIT Kerberos authentication is provided in HPSS 7.4; however, IBM Service Agreements for HPSS do not provide support for problem isolation nor fixing defects (Level 2 and Level 3 support) in MIT Kerberos. Kerberos maintenance and support must be site-provided.

Kerberos 5 is an option for the authentication mechanism. When this option is used, the local realm name is taken to be the name of a Kerberos realm. The Kerberos security services are used to obtain and verify credentials.

2.1.2.4. LDAP



LDAP authorization is not supported by IBM Service Agreements. The following information is provided for sites planning to use LDAP authorization with HPSS 7.4 as a site-supported feature.

An option for the authorization mechanism is to store HPSS security information in an LDAP directory. LDAP is a standard for providing directory services over a TCP/IP network. A server supporting the LDAP protocol provides a hierarchical view of a centralized repository of data and provides clients with sophisticated search options. The LDAP software supported by the HPSS LDAP authorization mechanism is IBM Tivoli Directory Server (Kerberos plug-in available for AIX only) and OpenLDAP (Kerberos plug-in available for AIX and Linux). One advantage of using the LDAP mechanism over the UNIX mechanism is that LDAP provides a central repository of information that is used by all HPSS nodes; it doesn't have to be manually kept in sync.

The rest of this section deals with how to accomplish various administrative tasks if the LDAP authorization mechanism is used.

Working with principals

Creating a principal

A principal is an entity with credentials, like a user or a server. The most straightforward way to create a new principal is to use the **-add** and **-ldap** options of the **hpssuser** utility. The utility will prompt for any needed information and will drive the **hpss_ldap_admin** utility to create a new principal entry in the LDAP server. To create a new principal directly with the **hpss_ldap_admin** utility, use the following command at the prompt:


```
princ create -uid <uid> -name <name> -gid <gid> -home <home> -shell  
<shell> [-uuid <uuid>]
```

If no UUID is supplied, one will be generated.

Deleting a principal

Likewise, use the **-del** and **-ldap** options of the **hpssuser** utility to delete the named principal from the LDAP server. To delete a named principal directly with the **hpss_ldap_admin** utility, use the following command at the prompt:

```
princ delete [-uid <uid>] [-name <name>] [-gid <gid>]  
[-uuid <uuid>]
```

You may supply any of the arguments listed. This command will delete any principal entries in the LDAP information that have the indicated attributes.

Working with groups

Creating a group

To create a new group, use the following command at the **hpss_ldap_admin** prompt:

```
group create -gid <gid> -name <name> [-uuid <uuid>]
```

If no UUID is supplied, one will be generated.

Deleting a group

To delete a group, use the following command at the **hpss_ldap_admin** prompt:

```
group delete [-gid <gid>] [-name <name>] [-uuid <uuid>]
```

You may supply any of the arguments listed. This command will delete any group entries in the LDAP information that have the indicated attributes.

Adding a member to a group

To add a principal to a group, use the following command at the **hpss_ldap_admin** prompt:

```
group add <principal> [-gid <gid>] [-name <name>] [-uuid <uuid>]
```

You may supply any of the arguments listed to select the group to which the named principal will be added.

Removing a member from a group

To remove a principal from a group, use the following command at the **hpss_ldap_admin** prompt:

```
group remove <principal> [-gid <gid>] [-name <name>] [-uuid <uuid>]
```

You may supply any of the arguments listed to select the group from which the named principal will be removed.

Working with trusted foreign realms

Creating a trusted foreign realm

To add an entry for a trusted foreign realm, use the following **hpss_ldap_admin** command:

```
trealm create -id <realmID> -mech <mechanism> -name <realmName> -url  
<url>
```

The arguments are as follows:

- *-id <realmID>* is the numeric realm ID for the foreign realm.
- *-mech <mechanism>* is a string identifying the authorization mechanism in use at the foreign realm, such as "unix" or "ldap".
- *-name <realmName>* is the name of the foreign realm; for example, "SOMEREALM.SOMEDOMAIN.COM".
- *-url <url>* is the URL of the security mechanism of the foreign realm. This only matters if the foreign realm is using LDAP as its authorization mechanism. If so, this must be the LDAP URL of the main entry for the security realm in the foreign LDAP server. This should be obtained from the foreign site's administrator. An example would be: "ldap://theirldapserver.foreign.com/cn=FOREIGNREALM.FOREIGN.COM".

Deleting a trusted foreign realm

To delete an entry for a trusted foreign realm, use the following **hpss_ldap_admin** command:

```
trealm delete [-id <realmID>] [-name <realmName>]
```

Any of the arguments listed can be supplied to select the trusted realm entry that will be deleted.

2.2. HPSS server security ACLs

Beginning with release 6.2, HPSS uses a table of access control information stored in the DB2 configuration database to control access to HPSS servers. This is the AUTHZACL table. HPSS software uses the configured authentication mechanism (for example, Kerberos) to determine a caller's identity via credentials provided by the caller, then uses the configured authorization mechanism to retrieve the details of the caller that determine the access granted. Once the identity and authorization information have been obtained, each HPSS server grants or denies the caller's request based on the access control list information stored in the database.

The default ACLs (Access Control Lists) for each type of server are as follows:

Core Server:

```
r--c--- user ${HPSS_PRINCIPAL_FTPD}
rw-c-dt user ${HPSS_PRINCIPAL_MPS}
rw-c-d- user ${HPSS_PRINCIPAL_SSM}
r--c--- user ${HPSS_PRINCIPAL_FS}
-----t any_other
```

Gatekeeper:

```
rw----- user ${HPSS_PRINCIPAL_CORE}
rw-c--- user ${HPSS_PRINCIPAL_SSM}
-----t any_other
```

Location Server:

```
r--c--t user ${HPSS_PRINCIPAL_SSM}
r-----t any_other
```

Mover:

```
r--c--t user ${HPSS_PRINCIPAL_PVL}  
rw-c--t user ${HPSS_PRINCIPAL_SSM}  
r-----t any_other
```

PVL:

```
rw-c-dt user ${HPSS_PRINCIPAL_CORE}  
rw--dt user ${HPSS_PRINCIPAL_PVR}  
rw-c-dt user ${HPSS_PRINCIPAL_SSM}  
-----t any_other
```

PVR:

```
rw-c-dt user ${HPSS_PRINCIPAL_PVL}  
rw-c--t user ${HPSS_PRINCIPAL_SSM}  
-----t any_other
```

RAIT Engine:

```
r--c--t user ${HPSS_PRINCIPAL_PVL}  
rw-c--t user ${HPSS_PRINCIPAL_SSM}  
r-----t any_other
```

SSM:

```
rwxcidt user ${HPSS_PRINCIPAL_ADM_USER}  
-----t any_other
```

All other types:

```
rw-c-dt user ${HPSS_PRINCIPAL_SSM}  
-----t any_other
```

In most cases, the ACLs created by default for new servers should be adequate. In normal operation, the only ACL that has to be altered is the one for the SSM client interface. This is handled automatically by the **-ssm** option of the **hpssuser** utility. If, for some reason, an ACL should need to be modified in some other way, the **hpss_server_acl** utility can be used. See the **hpss_server_acl** man page for more information.

2.3. SSM user security

SSM supports two types of users, administrators and operators:

administrator

This security level is normally assigned to an HPSS administrator. Administrators may open all SSM windows and perform all control functions provided by SSM.

operator

This security level is normally assigned to an HPSS operator. Operators may open most SSM windows and can perform all SSM control functions except for HPSS configuration.

Security is applied both at the window level and the field level. A user must have permission to open a window to do anything with it. If the user does succeed in opening a window, all items on that window may be viewed. Field level security then determines whether the user can modify fields, push buttons, or otherwise modify the window.

The security level of an SSM user is determined by his entry in the access control information table in the HPSS configuration database. The initial security level for a user is assigned when the SSM user is created by **hpssuser**. Security levels may be viewed and modified with the **hpss_server_acl** utility. See also *Section 3.3.2.2, “SSM user authorization”*.

2.4. Location policy

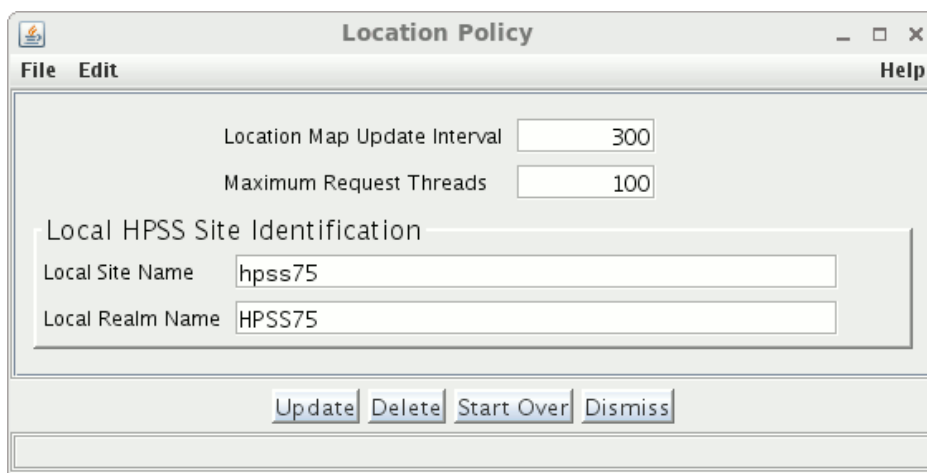
All Location Servers in an HPSS installation share a location policy. Only one location policy is permitted per HPSS installation. The location policy is used by the Location Servers to determine how and how often information should be updated. In general, most of the default values for the policy can be used without change.

2.4.1. Configuring and updating a location policy

The location policy can be created and updated using the *Location Policy* window. If the location policy does not exist, the fields will be displayed with default values for a new policy. Otherwise, the configured policy will be displayed.

Once a location policy is created or updated, it will not be in effect until all local Location Servers are restarted or reinitialized. The **Reinitialize** button on the *Servers* list window (see *Section 5.1, “Server list”*) can be used to reinitialize a running Location Server.

2.4.2. Location policy configuration window



Field descriptions.

Location Map Update Interval

Interval in seconds that the Location Server rereads the location map.



Advice

If this value is set too low, a load will be put on the database while reading configuration metadata. If set too high, new servers will not be registered in a timely manner. Set this value higher if timeouts are occurring during Location Server communications.

If you have multiple Location Servers, you should consider increasing the update interval since each Location Server obtains information independently and will increase the overall system load.

Maximum Request Threads

The maximum number of concurrent client requests allowed.



Advice

If the Location Server is reporting heavy loads, increase this number. If this number is above 300, consider replicating the Location Server on a different machine. Note if this value is changed, the general configuration thread value (Thread Pool Size) should be adjusted so that its value is always larger than the Maximum Request Threads. See *Section 5.2.1.3, “Interface controls”*.

Maximum Request Threads should not normally exceed (**Maximum Location Map Threads** + 400). This is not enforced. If you need more threads than this to handle the load, consider configuring an additional Location Server.

Local Site Name

The descriptive name of the HPSS installation.



Advice

Pick a name to uniquely describe the HPSS system.

Local Realm Name

The name where the realm containing Location Server path information is stored.



Advice

All clients need to know this group name since it is used by them when initializing to contact the Location Server. If the default is not used, ensure that the associated environment variable for this field is changed accordingly for all HPSS interfaces.

2.4.3. Deleting a location policy

The location policy can be deleted using the *Location Policy* window. Since only one location policy may be defined in a system and it must exist for the system to run, it is better to simply update the policy rather than delete and recreate it. See *Section 2.4.2, “Location policy configuration window”* for more details.

2.5. Restricting user access to HPSS

2.5.1. Overview

System administrators may deny access to various HPSS operations for a specific user by including that user in the Restrict Access File. This file is read by the Core Server at server startup. HPSS

system administrators may use the `magic` command line tool's `reinitrat` command to cause a running Core Server to reread the configuration file.



If an HPSS system administrator is restricted from all operations in the Restrict Access File, the system administrator's standard user ID will be restricted from using `magic`'s `reinitrat` command. The system administrator must run `magic` with root-equivalent privileges for this operation.

2.5.2. Configuration File

To set up the configuration file, add the `HPSS_RESTRICT_ACCESS_FILE` environment variable to `/var/hpss/etc/env.conf`. Set the value of this variable to the path name of the file that will contain the list of access restrictions. Ensure this environment variable is defined appropriately before starting the Core Server. To notify the Core Server of updates to the value of this environment variable, reinit the Core Server.

For example:

```
HPSS_RESTRICT_ACCESS_FILE=/var/hpss/etc/restrict-access-file
```



The file should be configured on the system where the root Core Server is running. Additionally, if running with multiple storage subsystems on different machines, be sure to configure the `HPSS_RESTRICTED_ACCESS_FILE` on each machine where a Core Server runs.

2.5.2.1. File Permissions

Examine the file system permissions for the Restrict Access File. You'll want to set the permissions to be permissive enough for the Core Server to read the file. And you'll want to set the permissions to be restrictive enough so that the file cannot be modified by people without authority to (un)restrict end-users' access to HPSS.

2.5.2.2. File Format

The Restrict Access File format is:

- one entry per user
- one user per line
- each line uses a CSV format that lists a user ID and one or more restrictions
- each line can end with a comment beginning with the `#` character

In general the file format looks like:

```
<UID> <restriction_1>, <restriction_2>, ..., <restriction_n> #comment
```



certain types of errors in a Restrict Access File can cause the Core Server to exit. These errors are typically misspelled restriction names or lines that are syntactically incorrect. In addition, finding a user ID in more than one line (more than one config file entry) will cause the Core Server to exit.

Please test any changes to the Restrict Access File prior to making the update on a production HPSS system. If an update to the Restrict Access File causes the Core Server to exit, revert to a previous version of the Restrict Access File and examine the updates for syntax errors.

2.5.3. ra_verify_config

The tool `ra_verify_config` is provided to, hopefully, detect and eliminate the types of errors mentioned above before these errors are discovered by the Core Server. `ra_verify_config` reads each line of the Restrict Access File and attempts to verify that the line is correct. It is expected that this verification will be done before submitting the file to the Core Server.

2.5.4. Restriction Types

Individual restrictions are listed for each user in the restrict access config file. The restriction names and meanings are:

- RestrictAll restrict the user from all actions
- RestrictFileCopies restrict the user from copying files
- RestrictFileCreates restrict the user from creating new files
- RestrictFileWrites restrict the user from writing to new and existing files
- RestrictStages restrict the user from staging files from lower levels (typically tape) to higher levels (typically disk) in storage hierarchies

2.5.5. Example Restrict Access File

```
4200 RestrictAll            # don't allow user ID 4200 to perform any HPSS operations
4042 RestrictFileCreates   # don't allow user ID 4042 to create any new files (write access)
4242 RestrictStages        # don't allow user ID 4242 to stage any data from tape
4420 RestrictFileCopies, RestrictFileWrites, RestrictFileCreates # don't allow user ID 4420 to perform any HPSS operations
```

Chapter 3. Using SSM

3.1. The SSM System Manager

3.1.1. Starting the SSM System Manager

Before starting the SSM System Manager (SM), review the SM key environment variables described in the *Storage System Management* section of the *HPSS Installation Guide* (in the *HPSS planning* chapter). If the default values are not desired, override them using the **hpss_set_env** utility. See the **hpss_set_env** man page for more information.

To start the SM, invoke the **rc.hpss** script as follows:

```
% su -  
% /opt/hpss/bin/rc.hpss -m start
```

3.1.2. Tuning the System Manager RPC thread pool and request queue sizes

Tuning the System Manager RPC thread pool and request queue sizes can improve the performance of both the System Manager and its clients (**hpssgui** and **hpssadm**). It is not necessary, however, to do the tuning when bringing up SSM for the first time. In fact, it can be helpful to postpone the tuning until after the site has a chance to learn its own SSM usage patterns.

The System Manager client interface RPC thread pool size is defined in the **Thread Pool Size** field on the Interface Controls tab of the System Manager's *Core Server Configuration* window (see Section 5.2.1.3, "Interface controls"). This is the maximum number of RPCs that can be active at any one time for the client interface (that is, all the **hpssgui** and **hpssadm** clients). For the server RPC interface (connections to the SSM System Manager from other HPSS servers), this value is determined by the `HPSS_SM_SRV_TPOOL_SIZE` environment variable.

The System Manager client interface RPC request queue size is defined in the **Request Queue Size** field on the Interface Controls tab of the System Manager's *Core Server Configuration* window (see Section 5.2.1.3, "Interface controls"). This is the maximum number of RPC requests from **hpssgui** and **hpssadm** clients which can be queued and waiting to become active. For the server RPC interface, this value is determined by the `HPSS_SM_SRV_QUEUE_SIZE` environment variable.

Ideally, if the site runs many clients, the client interface RPC thread pool size should be as large as possible; the default is 100. Testing this value at 300 showed the System Manager memory size more than doubled. The larger RPC thread pool size makes the System Manager much more responsive to the many clients but also requires it to use more memory.

Experimentation shows that leaving the client interface RPC thread pool size at 100 and leaving the client interface RPC request queue size at its default (600) works pretty well for up to about 40 clients. During further experimentation, setting the client interface RPC request queue size to 1000 resulted in

very little effect on memory usage; with 40 clients connected, the client interface RPC request queue used never went above about 500, but the client interface RPC thread pool was constantly filled.

Avoid allowing the client interface RPC thread pool to become full. When this happens, new RPCs will be put into the client interface RPC request queue to wait for a free thread from the thread pool. This makes the client response appear slow, because each RPC request is having to wait its turn in the queue. To help mitigate this, when the thread pool is full, the System Manager notifies all the threads in the thread pool that are waiting on list updates to return to the client as if they just timed out as normal. This could be as many as 15 threads per client that are awakened and told to return, which makes those threads free to do other work.

If the client interface RPC thread pool is still full (as it could be if, for example, there were 15 threads in the client interface RPC request queue that took over the 15 that were just released), then the System Manager sets the wait time for the new RPCs to one second rather than whatever the client requested. This way the RPC won't try to hang around too long.

Realize that once the System Manager gets in this mode (constantly having a full client interface RPC thread pool and having to cut short the thread wait times), the System Manager starts working hard and the CPU usage will start to increase. If you close some windows, some clients, or some of both, things should start to stabilize again.

You can see whether the System Manager client interface RPC thread pool has ever been full by looking at the **Maximum Active/Queued RPCs** field in the Client column of the RPC Interface Information group in the *System Manager Statistics* window (see *Section 3.9.4.1, "System Manager Statistics window"*). If this number is greater than or equal to the corresponding client interface's Thread Pool Size (default 100), then the thread pool was full at some time during the System Manager execution (although it may not be full currently).

To tell whether the thread pool is currently full, look at the number of Queued RPCs. If Queued RPCs is zero, then the thread pool is not full at the moment.

If Active RPCs is equal to Thread Pool Size, then the thread pool for the interface is currently full. Active RPCs should never be greater than Thread Pool Size. When it reaches the Thread Pool Size, then the new RPCs will be queued and Queued RPCs become greater than zero.

When the thread pool gets full, the System Manager tries harder to clear them out before accepting new ones, to ensure that if the thread pool fills up, it doesn't stay full for long.

If the site runs with low refresh rates and more than 40 clients, the recommendation is to set the client interface RPC thread pool size to 150 or 200 and the client interface RPC request queue size to 1000 in the System Manager *Server Configuration* window (see *Section 5.2.1.3, "Interface controls"*). Otherwise, the default values should work well.

3.1.3. Labeling the System Manager RPC program number

Labeling the System Manager RPC program number is not required but can be a useful debugging aid.

The SSM System Manager registers with the RPC portmapper at initialization. As part of this registration, it tells the portmapper its RPC program number. Each HPSS server configuration

contains the server's RPC program number. To find the System Manager's program number, open the *Servers* window, select the SSM System Manager, and click the **Configure** button to open the *SSM System Manager Configuration* window. The System Manager's RPC program number is in the **Program Number** field on the Execution Controls tab of this window.

The **rpcinfo** utility with the **-p** option will list all registered programs, their RPC program numbers, and the port on which they are currently listening for RPCs. When diagnosing SSM problems, it can be useful to run the **rpcinfo** program and search for the System Manager RPC program number in the output. Look to see whether the System Manager has successfully initialized its RPC interface and to see which port **hpssgui** and **hpssadm** clients must access to reach the System Manager.

This task can be made a bit easier if the System Manager RPC program number is labeled in the portmapper. To do this, add a line for the System Manager in the `/etc/rpc` file specifying the program number and a convenient RPC service name such as "hpss_ssm" (note that names may not contain embedded spaces). Then this service name will show up in the **rpcinfo** output.

The format of the `/etc/rpc` file differs slightly across platforms. See the platform-specific man pages for the `rpc` file for details. The **rpcinfo** utility is typically found in either `/usr/bin` (AIX) or `/usr/sbin` (Linux).

3.2. Quick startup of hpssgui

We recommend that **hpssgui** sessions be invoked from the user's desktop computer instead of on the HPSS server machine. **hpssgui** is an application designed to run in the Java environment on the user's desktop computer and to communicate with the remote SSM System Manager. If **hpssgui** is executed on the remote System Manager host, it must run through an X Window session and it may run very slowly in that environment. This is a limitation of Java and networks.

We recognize the value of using the remote X functionality as a quick way to get SSM running, but once your system is up, it is highly recommended that you configure local desktop SSM **hpssgui** clients for all HPSS administrators and operators. Local desktop **hpssgui** configuration is detailed in *Section 3.3, "Configuration and startup of hpssgui and hpssadm"* below.

The following are steps for quickly configuring and starting an SSM GUI client:

1. Use the **hpssuser** utility to create an SSM user with admin authority. See *Section 3.3.2.1, "The hpssuser utility"* and the **hpssuser** man page for more information.
2. The **mkhpss** utility generates the `ssm.conf` SSM configuration text file when configuring the SM. See the *Configure HPSS infrastructure* section of the *HPSS Installation Guide* (in the *HPSS installation and infrastructure configuration* chapter), for more details. Verify the existence of the file `$HPSS_PATH_SSM/ssm.conf`.
3. Start the **hpssgui** script:

```
% /opt/hpss/bin/hpssgui.pl
```

- Note that the **-m** option can be used to specify the desired SSM configuration file to be used. When this option is not specified, `hpssgui.pl` looks for the `ssm.conf` configuration file in the current directory, then in the directory defined by the `HPSS_PATH_SSM` environment variable

(usually `/var/hpss/ssm`). If the script doesn't find a configuration file in either directory, it will use default values to start the client.

- Note that the **-d** (debug) and **-S** (log file name) options can be used to capture all levels of **hpssgui** logging in a text file. Bear in mind, however, that this can generate significant amounts of log data. (See the **hpssgui** man page.)
- When you have decided on the **hpssgui** command line that is best for your installation, it will probably be useful to put the command in a shell script for the convenience of all SSM administrators and operators. For example, create a file called "gui" and put the following in it:

```
/opt/hpss/bin/hpssgui.pl \  
-m /my_directory/my_ssm.conf \  
-d \  
-S /tmp/hpssguiSessionLog.$(whoami)
```

Refer to the **hpssgui** man page for an extensive list of command line options. For example, some sites prefer to set the date format to a USA military format using the **-D "kk:mm:ss dd-MMM-yyyy"** option. Additionally, *Section 3.3.3, "SSM configuration file"* below provides a table of variables you can set in the SSM configuration file instead of using command line options; this section also covers all the various files that the **hpssgui** script uses.

3.3. Configuration and startup of hpssgui and hpssadm

This section describes in detail the procedures for configuring SSM and creating an SSM user account with the proper permissions to start up an **hpssgui** or **hpssadm** session. It also explains how to install the SSM client on the user's desktop (the recommended configuration for **hpssgui**) and how to deal with special situations such as firewalls.

In the discussion that follows, *authentication* ensures that a user is who they claim to be relative to the system. *Authorization* defines the user's rights and permissions within the system.

Like other components of HPSS, SSM authenticates its users by using either Kerberos or UNIX. Users of the **hpssgui** and **hpssadm** utilities are authenticated to SSM by either a Kerberos principal and a password or by a UNIX username and a password. The System Manager must be configured to use the appropriate authentication and a Kerberos principal or UNIX user account must be created for each SSM user.

Unlike other components of HPSS, SSM does not use LDAP or UNIX to authorize its users. SSM users are authenticated based on their entries in the HPSS DB2 AUTHZACL table. Through this table, SSM supports two levels of SSM client authorization:

admin	This security level is normally assigned to an HPSS administrator. The admin user can view all SSM windows and perform all control functions provided by SSM.
operator	This security level is normally assigned to an HPSS operator. The operator user can view most SSM windows and perform all SSM control functions except for changing the HPSS configuration.

Configuration of an SSM user requires that:

1. The System Manager is configured to accept the desired authentication mechanism.
2. The proper user accounts are created:
 - UNIX or Kerberos accounts are created for the user authentication.
 - The proper authorization entries for the user are created in the AUTHZACL table.
3. The proper SSM configuration files are created and installed.

See *Section 3.3.1, “Configuring the System Manager authentication for SSM clients”*, *Section 3.3.2, “Creating the SSM user accounts”*, and *Section 3.3.3, “SSM configuration file”* for the procedures for these tasks.

See *Section 3.3.4, “SSM help files (optional)”* for instructions on installing the SSM help package.

See *Section 3.3.5, “SSM desktop client packaging”* for instructions for installing **hpssgui** or **hpssadm** on the user’s desktop.

See *Section 3.3.6, “Using SSM through a firewall”* for advice about using **hpssgui** or **hpssadm** through a network firewall.

3.3.1. Configuring the System Manager authentication for SSM clients

The System Manager is configured initially by **mkhpss** for new HPSS systems or by the conversion utilities for upgraded HPSS systems to use the proper authentication mechanism.

If it is necessary later to modify the authentication mechanism for **hpssgui** or **hpssadm** users, or to add an additional mechanism, bring up the *Servers* window, select the System Manager, and click the **Configure** button. On the *System Manager Configuration* window, select the Interface Controls tab. For the SSM Client Interface, make certain the check box for the desired authentication mechanism, KRB5 or UNIX, is selected. Both mechanisms may be enabled if desired.

Next, select the Security Controls tab. Make certain one of the authentication service configurations is set to use the desired **Mechanism**, an **Authenticator Type** of Keytab, and a valid keytab file name for **Authenticator** (default is `/var/hpss/etc/hpss.keytab` for Kerberos, or `/var/hpss/etc/hpss.unix.keytab` for UNIX).

To remove an authentication mechanism from the System Manager, so that no SSM user may be authenticated using that mechanism, reverse the above process. Unselect the mechanism to be removed from the SSM Client Interface on the Interface Controls tab. On the Security Controls tab, change the **Mechanism** and **Authenticator Type** fields of the mechanism to be removed to Not Configured, and change **Authenticator** to blank.

See *Section 5.2.1.3, “Interface controls”*, and *Section 5.2.1.4, “Security controls”* for more information.

3.3.2. Creating the SSM user accounts

3.3.2.1. The hpssuser utility

The **hpssuser** utility is the preferred method for creating, modifying or deleting SSM users. It creates the necessary UNIX or Kerberos accounts. It creates an entry in the AUTHZACL table for the user with the proper authorization.

The following is an example of using the **hpssuser** utility to provide administrative access to SSM to user "john". In this example, the user already has either a UNIX or Kerberos account.

```
% /opt/hpss/bin/hpssuser -add john -ssm
[ adding ssm user ]
1) admin
2) operator
Choose SSM security level
(type a number or RETURN to cancel):
> 1
[ ssm user added : admin ]
```

After SSM users are added, removed, or modified, the System Manager will automatically discover the change when the user attempts to log in. See the **hpssuser** man page for details.



Removing an SSM user or modifying an SSM user's security level won't take effect until that user attempts to start a new session. This means that if an SSM user is removed, any existing SSM sessions for that user will continue to work; access won't be denied until the SSM user attempts to start a new SSM session. Likewise, if the SSM user's security level is changed, any existing sessions for that user will continue to work at the old security level; the new security level access won't be recognized until the SSM user starts a new SSM session.

3.3.2.2. SSM user authorization

SSM user authorization is set properly by the **hpssuser** utility with no further modification required. This section explains how the authorization levels are stored internally and how they may be viewed for debugging or modified.

The SSM admin and operator security authorization levels are defined in the AUTHZACL table in the HPSS DB2 database. Each SSM user must have an entry in this table. The permissions supported in the table are:

- **r** indicating *read*
- **w** indicating *write*
- **x** indicating *execute*
- **c** indicating *control*
- **i** indicating *insert*

- **d** indicating *delete*
- **t** indicating *test*

SSM administrators must be granted all permissions: `rwxcidt`. SSM operators must be granted `r-c-t` permissions. All other permission combinations are not recognized by the SSM server and will be treated as no permissions at all.

The AUTHZACL table may be viewed or updated with the **hpss_server_acl** utility. The **hpssuser** utility program creates and deletes SSM user entries in the AUTHZACL table using the **hpss_server_acl** utility. Normally, there is no need to invoke the **hpss_server_acl** utility directly because it is invoked by the **hpssuser** utility. However, it is a useful tool for examining and modifying the authorization table.



*Access to the **hpss_server_acl** program, **hpssuser** program, HPSS DB2 database, and all HPSS utility programs should be closely guarded. An operator with permission to run these tools could modify the type of authority granted to anyone by SSM. Note that access to the database by many of these tools is controlled by the permissions on the `/var/hpss/etc/mm.keytab` file.*

Here is an example of using the **hpss_server_acl** utility to set up a client's permissions to be used when communicating with the SSM server. Note that the **default** command should be used only when creating the ACL for the first time, as it removes any previous entries for that server and resets all the server's entries to the default values:

```
% /opt/hpss/bin/hpss_server_acl
hsa> acl -t SSM -T ssmclient
hsa> show
hsa> default # Note: ONLY if creating acl for the first time
hsa> add user <username> <permissions>
hsa> show
hsa> quit
```

If the ACL already exists, this command sequence gives user "bill" operator access:

```
% /opt/hpss/bin/hpss_server_acl
hsa> acl -t SSM -T ssmclient
hsa> show
hsa> add user bill r--c--t
hsa> show
hsa> quit
```



Removing an SSM user or modifying an SSM user's security level won't take effect until that user attempts to start a new session. This means that if an SSM user is removed, any existing SSM sessions for that user will continue to work; access won't be denied until the SSM user attempts to start a new SSM session. Likewise, if the SSM user's security level is changed, any existing sessions for that user will continue to work at the old security level; the new security level access won't be recognized until the SSM user starts a new SSM session.

3.3.2.3. User keytabs

A keytab is a file containing a user name and an encrypted password. The keytab file can be used by a utility program to perform authentication without human interaction or the need to store a password

in plain text. Each user who will run the **hpssadm** or **hpssgui** utility may employ a keytab. It is recommended that one keytab file per user be created rather than one keytab containing multiple users.

Each keytab file should be readable only by the user for whom it was created. Each host from which the **hpssadm** or **hpssgui** utility is executed must be secure enough to ensure that the user's keytab file cannot be compromised. An illicit process which gained access to a Kerberos keytab file could gain the user's credentials anywhere in the Kerberos realm; one which gained access to a UNIX keytab file could gain the user's credentials at least on the System Manager host. The **hpssgui** and **hpssadm** utilities will look for a default keytab in the user's home directory, called `krb5.keytab` for Kerberos or `keytab` for UNIX.

Keytabs are created for the user by the **hpssuser** utility when the keytab file path is specified. Keytabs may also be created manually with the **hpss_krb5_keytab** or **hpss_unix_keytab** utility, as described below.

Keytabs for Kerberos authentication: **hpss_krb5_keytab**.

The **hpss_krb5_keytab** utility may be used to generate a keytab with Kerberos authentication in the form usable by the **hpssadm** and **hpssgui** programs. See the **hpss_krb5_keytab** man page for details.

The Kerberos keytab is interpreted by the KDC of the Kerberos realm specified by the **hpssadm** and **hpssgui** utilities (see the **-k** and **-u** options on the respective man pages). This must be the same Kerberos realm as that used by the System Manager. This means the **hpss_krb5_keytab** utility must be executed on a host in the same realm as the System Manager.

This example for a user named *joe* creates a Kerberos keytab file named `krb5.keytab`:

```
% /opt/hpss/bin/hpss_krb5_keytab -f /home/joe/krb5.keytab
HPSS_ROOT is not set; using /opt/hpss
KRB5_INSTALL_PATH is not set; using /krb5
password:
Your keytab is stored at /home/joe/krb5.keytab
```

Keytabs for UNIX authentication: **hpss_unix_keytab**.

The **hpss_unix_keytab** utility may be used to generate a keytab with UNIX authentication in the form usable by the **hpssadm** and **hpssgui** programs. See the **hpss_unix_keytab** man page for details.

The UNIX keytab is interpreted on the host on which the System Manager runs, not the host on which the **hpssadm** client utility runs. The encrypted password in the keytab must match the encrypted password in the password file on the System Manager host. Therefore, the **hpss_unix_keytab** utility must be executed on the host on which the System Manager runs.

The **hpss_unix_keytab** utility must be able to read the user's encrypted password from the password file.

This example for a user named *joe* creates a UNIX keytab file named `keytab`:

```
% /opt/hpss/bin/hpss_unix_keytab -f /home/joe/keytab add joe
```

This command copies the encrypted password from the password file into the keytab.

Note: The **-r** option of the **hpss_unix_keytab** utility will first randomize the user's password and then place it into the keytab file. This is rarely desirable. The **-p** option will change the password first, and then place it into the keytab file. If multiple keytab files exist for this user, this option will invalidate the other keytabs for this user, even if the password is the same. Therefore, this option is also rarely desirable.

3.3.3. SSM configuration file

The **hpssgui** and **hpssadm** scripts use the SSM configuration file `ssm.conf` for configuration.

The **mkhpss** utility creates the SSM configuration file for the security mechanism supported by SSM. The **mkhpss** utility stores the generated `ssm.conf` at `$HPSS_PATH_SSM`; the default location is `/var/hpss/ssm`. The configuration file contains host-specific and site-specific variables that the **hpssgui** and **hpssadm** script reads. The variables contain information about:

- SSM hostname
- SSM RPC number
- SSM RPC protection level
- SSM security mechanism
- SSM UNIX realm (only if using UNIX authentication)

If any of these configuration parameters are modified, the `ssm.conf` file must be updated and redistributed from the server machine to all of the SSM client machines.

Users can also use their SSM configuration file to manage SSM client parameters instead of using the command line options. The **hpssgui** and **hpssadm** scripts can be directed to use an alternate SSM configuration file with the **-m** option. The default SSM configuration file contains comments describing each of the available parameters that may be set along with any associated environment variable and command line option. The following table documents these variables and the corresponding command line options:

Table 3.1. SSM general options

File Option	Command Line Option	Functionality
HPSS_SSM_ALARM_RATE	-A	Alarm refresh rate
LOGIN_CONFIG	-C	Full path to <code>login.conf</code> file
HPSS_SSM_DATE_FORMAT	-D	Date format pattern
HPSS_SSM_ALARMS_GET	-G	Number of alarms requested per poll
HPSS_SSM_LIST_RATE	-L	How long hpssgui / hpssadm waits between polling for lists
HPSS_SSM_ALARMS_DISPLAY	-N	Max number of alarms displayed by hpssgui
HPSS_SSM_CLASSPATH	-P	Full path to <code>hpss.jar</code> file

File Option	Command Line Option	Functionality
LOG_FILE	-S	Full path for session log file
HPSS_SSM_WAIT_TIME	-W	How long SM waits before returning if object is unchanged
HPSS_SSM_CNTRY_CODE	-c	Country code, internationalization
HPSS_SSM_DEBUG	-d	Debug flag
HPSS_SSM_SM_HOST_NAME	-h	System manager hostname
HPSS_SSM_USER_PREF_PATH	-i	Path to SSM preferences
JAVA_BIN	-j	Path to Java bin directory
KRB5_CONFIG	-k	Full path to <code>krb5.conf</code> file
HPSS_SSM_LANG_CODE	-l	Language code, internationalization
SSM_CONFIG	-m	Full path to SSM configuration file
HPSS_SSM_SM_PORT_NUM	-n	port number <i>or</i> RPC number:program number
HPSS_SSM_CLIENT_IP	-p	Client IP address
HPSS_SSM_RPC_PROT_LEVEL	-r	RPC protection level
HPSS_SSM_SEC_MECH	-s	Security mechanism
HPSS_SSM_UNIX_REALM	-u	UNIX realm

Table 3.2. HPSSGUI-specific options

File Option	Command Line Option	Functionality
HPSSGUI_AUTHENTICATOR	-a	Authenticator (keytab path name)
HPSSGUI_LOOK_AND_FEEL	-F	Look and feel
HPSSGUI_MO_RATE	-M	How long hpssgui waits between polling for managed objects
HPSSGUI_METAL_THEME	-T	Theme file, for look and feel
HPSSGUI_METAL_BG	-b	Background color
HPSS_SSM_HELP_FILES_PATH	-f	Path to help files
HPSS_SSM_HELP_URL_TYPE	-g	Help path URL type

Table 3.3. HPSSADM-specific options

File Option	Command Line Option	Functionality
HPSSADM_USER_NAME	-U	User name for hpssadm
HPSSADM_AUTHENTICATOR	-a	Authenticator (keytab path name)
HPSSADM_BATCH_MODE	-b	Batch mode flag

Information on tuning client polling rates for optimal performance is available in the **hpssadm** and **hpssgui** man pages.

Options are specified, in precedence order, by: 1) the command line; 2) the user's environment (see the man pages for environment variable names); 3) the SSM configuration file; or 4) internal default values.

3.3.3.1. login.conf

The `login.conf` file is a login configuration file that specifies the security authentication required for the **hpssgui** and **hpssadm** programs. A copy of the `login.conf` file is included in the `hpss.jar` file and should require no site customization. However, a template for the file is provided in `/opt/hpss/config/templates/login.conf.template` should the site need to customize the security mechanisms.

See the `/opt/hpss/config/templates/login.conf.template` file for details.

3.3.3.2. krb5.conf (for use with Kerberos authentication only)

The `krb5.conf` file is the Kerberos configuration file which allows the client to authenticate to the Kerberos realm. This file is only required if Kerberos authentication is used. The Kerberos installation process generates a default Kerberos configuration file in `/etc/krb5.conf`.

The following is an example of this file. Realm names, hostnames, and the other data must be customized to operate properly in the user's site environment.

krb5.conf:

```
[logging]
default = FILE:/var/hpss/log/krb5libs.log
kdc = FILE:/var/hpss/log/krb5kdc.log
admin_server = FILE:/var/hpss/log/kadmind.log

[libdefaults]
ticket_lifetime = 24000
default_realm = EXAMPLE.COM
default_keytab_name = /etc/v5srvtab
default_tkt_etypes = aes256-cts-hmac-shal-96 aes128-cts-hmac-shal-96
des3-cbc-shal arcfour-hmac-md5
default_tgs_etypes = aes256-cts-hmac-shal-96 aes128-cts-hmac-shal-96
des3-cbc-shal arcfour-hmac-md5

[realms]
EXAMPLE.COM = {
kdc = example.com:88
admin_server = example.com:749
}

[domain_realm]
example.com = EXAMPLE.COM
```

Note that having strong encryption types can cause authentication failures. Newer Kerberos libraries may define new encryption types that older libraries don't know of. Strong encryption types are known to fail for Java implementations that have the standard JCE policy files installed.

3.3.4. SSM help files (optional)

The SSM help files are an HTML version of the *HPSS Management Guide*. Individual sections of this guide are available from the Help menus on the SSM windows.

To access help windows from the **hpssgui**, the help files must be accessible from each client machine. We recommend storing these files in a file system shared by the clients so that they don't need to be installed on every SSM client machine. By default, the **hpssgui** script looks for the help files in `$HPSS_HELP_FILES_PATH`. The default location is `/hpss_src/hpss-<version>/ssmhelp` and can be overridden by using the **-f** option.

Help files are distributed with HPSS or can be downloaded from the HPSS web site. They should be installed in the `$HPSS_HELP_FILES_PATH` location and/or the path specified by the **-f** option. Refer to the *Locate HPSS documentation and set up manual pages* subsection of the *HPSS Installation Guide* (in the *HPSS installation and infrastructure configuration* chapter), for instructions on how to install the help files. See the **hpssgui** man page for more details.

3.3.5. SSM desktop client packaging

A full installation of HPSS is not needed on machines used only for executing **hpssgui** or **hpssadm**. These machines, referred to here as "SSM client machines", only require the proper version of Java plus a subset of HPSS components.

It is strongly recommended that a desktop configuration be created and installed for each **hpssgui** user. The **hpssgui** program may run very slowly if it is executed on the System Manager machine and displayed back to the user's desktop via remote X.

There is no advantage to executing the **hpssadm** program on the desktop machine. It will perform just as well when executed remotely as on the desktop. In fact, it is recommended that **hpssadm** be executed on the System Manager machine rather than on the user's desktop since this simplifies the dissemination and protection of the user keytabs. Nonetheless, instructions are included here for packaging the **hpssadm** for sites who have a need to execute it on the desktop.

If the SSM code on the System Manager machine is recompiled, or if the System Manager is reconfigured, the client package will no longer work since it is then possible the `hpss.jar` file may be out of sync with the System Manager code. **Since each client will have its own copy of the `hpss.jar` file, the `hpss.jar` file should be redistributed to each client. This can be done by redistributing the entire SSM client package or by just redistributing the `hpss.jar` file.**

Section 3.3.5.1, "Automatic SSM client packaging and installation" describes how to use the **hpssuser** utility to package these components. *Section 3.3.5.2, "Manual SSM client packaging and installation"* describes how to select and package the components manually.

3.3.5.1. Automatic SSM client packaging and installation

The **hpssuser** utility provides a mechanism for packaging all the necessary client files required to execute the **hpssgui** program on the user's desktop host. Refer to the **hpssuser** man page for more information on generating an SSM Client Package. These files may also be copied manually; see *Section 3.3.5.2, "Manual SSM client packaging and installation"* for a list of the required files.

This example creates an SSM Client Package named "ssmclient.tar":

```
%/opt/hpss/bin/hpssuser -ssmclientpkg ssmclient.tar
[ packaging ssm client ]
[ creating ssmclient.tar ]
hpssgui.pl
hpssgui.vbs
hpss.jar
krb5.conf
ssm.conf
[ packaged ssm client in ssmclient.tar ]
```

Once the SSM Client Package has been generated simply FTP the tar file over to the client node and then extract the member files to the desired location.

3.3.5.2. Manual SSM client packaging and installation

This section describes the manual installation of the necessary client files required to execute the **hpssgui** or **hpssadm** program on the user's desktop host. The **hpssuser** utility also provides a mechanism for packaging these files automatically; see *Section 3.3.5.1, "Automatic SSM client packaging and installation"*.

The desktop machine requires the proper version of Java and the following HPSS files, which should be copied from the host on which the SSM System Manager executes:

- **scripts:** `hpssgui.pl`, `hpssgui.vbs`, `hpssadm.pl`, or `hpssadm.vbs`
- `hpss.jar`
- `ssm.conf`
- `krb5.conf` (if using Kerberos authentication)
- user keytab (if using **hpssadm**)
- help files (optional)

These are the default locations of these files on the SSM System Manager host, from which they may be copied:

startup scripts	/opt/hpss/bin
hpss.jar	/opt/hpss/bin
ssm.conf	/var/hpss/ssm
krb5.conf	/etc/krb5.conf
keytab file	/var/hpss/ssm/keytab.USERNAME
help files	/hpss_src/hpss-<version>/ssmhelp

These files may be installed in any location on the SSM client machines. The user must have at least read access to the files.

The SSM startup scripts `hpssgui.pl`, `hpssgui.vbs`, `hpssadm.pl`, and `hpssadm.vbs` provide the user with a command line mechanism for starting the SSM client. The `hpssgui.pl` script is a Perl script for starting the SSM Graphical User Interface and the `hpssadm.pl` script is a Perl script for starting the SSM Command Line User Interface. These scripts work on AIX, Linux, or Windows platforms

so long as Perl is installed on the host. The `hpssgui.vbs` script is a Visual Basic script for starting the Graphical User Interface and the `hpssadm.vbs` script is a Visual Basic script for starting the SSM Command Line User Interface. These scripts work only on Windows platforms.

These scripts depend on the ability to read the other files in the package. See the **hpssgui** and **hpssadm** man pages for details.

The `hpss.jar` file contains the **hpssadm** and **hpssgui** program files. This is stored on the server machine under `$HPSS_PATH_BIN`; the default location is `/opt/hpss/bin`. If the SSM source code on the server machine is recompiled or upgraded, the `hpss.jar` file must be redistributed to all of the SSM client machines.

The keytab is used only by the **hpssadm** program. See *Section 3.3.2.3, “User keytabs”* for details.

See *Section 3.3.4, “SSM help files (optional)”* for a description of the Help Files.

A writable directory is required for **hpssgui** or **hpssadm** session logs, if these are desired. The session log is an ASCII file that stores messages generated by the **hpssadm** or **hpssgui** programs. By default, the **hpssgui** and **hpssadm** scripts do not create session logs, but it is strongly encouraged that this capability be enabled by using the **-S <location>** option when running the script. The recommended location is `/tmp` on UNIX-like systems or `c:\tmp` on Windows systems. See the **hpssgui** and **hpssadm** man pages for more information on creating a session log. Having the session log available helps when debugging problems with the SSM client applications. It is the first thing that the SSM developers will ask for when someone is having problems with the **hpssgui** or **hpssadm**.

3.3.6. Using SSM through a firewall

3.3.6.1. The firewall problem

hpssgui and **hpssadm** require the use of several network ports which may be blocked if the client and System Manager are on opposite sides of a network firewall. Up to three ports may be affected:

- **hpssgui** and **hpssadm** must be able to access the port upon which the System Manager listens for requests.
- If the System Manager follows the default behavior of letting the portmapper select this port, then **hpssgui** and **hpssadm** also need access to port 111 in order to ask the portmapper where the System Manager is listening.
- If Kerberos authentication is used, then **hpssgui** and **hpssadm** additionally need access to port 88 to communicate with the KDC.

3.3.6.2. Solutions for operating through a firewall

SSM can operate through a firewall in three different ways:

- The **hpssgui** and **hpssadm** can use ports exempted by the network administrator as firewall exceptions. See the **-n** option described in the **hpssgui** and **hpssadm** man pages.
- The **hpssgui** and **hpssadm** can contact the System Manager across a Virtual Private Network (VPN) connection. See the **-p** and **-h** options described in the **hpssgui** and **hpssadm** man pages.

- The **hpssgui** and **hpssadm** can contact the System Manager across an SSH tunnel. See the instructions for tunneling in the **hpssgui** man page.

The firewall exception is the simplest of these. However, security organizations are not always willing to grant exceptions.

The VPN option is usually simple, regardless of how many ports are needed, but requires the site to support VPN. The site must also allow the VPN users access to the ports listed in *Section 3.3.6.1, “The firewall problem”*; not all sites do.

The SSH tunneling option has the advantage that it can be used almost anywhere at no cost. It has the disadvantage that the tunnel essentially creates its own firewall exception. Some security organizations would rather know about any applications coming through the firewall and what ports they are using rather than have users create exceptions themselves without the awareness of security personnel. A second disadvantage of tunneling is that if a particular client machine is compromised, any tunnels open on that client could also be compromised. The client machine may become a point of vulnerability and access to the other machines behind the firewall. A third disadvantage is that tunneling can be complex to set up, requiring slight or significant variations at every site.

The firewall and tunneling options both benefit from reducing the number of ports required:

- The need for port 111 can be eliminated by making the System Manager listen on a fixed port. To do this, set the `HPSS_SSM_SERVER_LISTEN_PORT` variable in the `env.conf` file to the desired port and restart the System Manager. Then use the `-n` option with the **hpssgui** and **hpssadm** startup scripts to specify this port.
- The need for port 88 can be eliminated only by avoiding Kerberos and using UNIX authentication.
- There is no way to eliminate the need for the port on which the System Manager listens.

3.3.6.3. Example: Using hpssgui through a firewall

Here is an example of how a particular site set up their **hpssgui** SSM client sessions using `krb5` authentication outside a firewall. Many of the items are site-specific, so modifications will need to be made to suit each site’s specific needs. Where this procedure would differ for a site using UNIX authentication, the UNIX instructions are also included.

At this site, VPN users were not allowed access to all the ports listed in *Section 3.3.6.1, “The firewall problem”* so they had to use a combination of VPN and SSH tunneling.

- Create a directory on the client machine to hold the SSM client files. It is recommended that a separate directory be created for each server hostname that the client will contact.
- Verify that the proper version of Java is installed. Add the Java `bin` directory to the user’s `$PATH`, or use the `-j` switch in the **hpssgui** script, or set `JAVA_BIN` in the user’s `ssm.conf` file. Java can be downloaded from <http://www.java.com>.
- Obtain files from the server machine:
 - Obtain the preferred **hpssgui** script for the client system from `/opt/hpss/bin` on the server machine and place it in the directory created on the client machine (see *Section 3.3.5, “SSM*

desktop client packaging”). There are several script options. Only one version of the script is needed:

- **hpssgui.pl** which is written in Perl and can be used on any system that has Perl installed. This is true for any major UNIX operating systems as well as Mac OS. For Windows users, Perl must be installed to use this version of the script. Users can easily obtain this from the web. A good Perl distribution for Windows is available at <http://www.activestate.com>.
- **hpssgui.vbs** is a Visual Basic Script version for Windows users. This version requires no prerequisite software.
- Obtain the **ssm.conf** file from `/var/hpss/ssm` on the server machine and place it in the directory where the **hpssgui** script resides. Alternately, specify the file to the **hpssgui** script with the **-m** option, if desired.
- Obtain the **hpss.jar** file from `/opt/hpss/bin` on the server machine and place it in the directory where the **hpssgui** script resides. If FTP is used to copy the file, make sure the copy is done in binary mode. If the file is installed in a different directory, specify it to the **hpssgui** script with the **-P** option, or by using configuration file settings or the appropriate environment variable (see the **hpssgui** man page).
- If Kerberos authentication is used, be sure to get the **krb5.conf** file that resides on the SSM server. This file should be located at `/etc/krb5.conf`. Place this file on the client machine in the directory where the **hpssgui** script resides. Alternately, specify this file to the **hpssgui** script with the **-k** option. Verify that UDP port 88 on the SSM Server machine is accessible; if not, then **hpssgui** will fail.
- To get access to ports inside the firewall, we can use a VPN connection or one or more SSH tunnels.
- Using a VPN connection will make it appear that we are inside the firewall. In this case, no tunnels are needed. If the firewall does not permit SSH connections, SSH tunnels cannot be used. Set up the VPN connection on the client machine.
- If using one or more SSH tunnels is preferred, on the SSM server machine, set the **HPSS_SSM_SERVER_LISTEN_PORT** variable in the **env.conf** file to a specific port (for example, 49999). Restart the System Manager so that it will recognize this variable.

On the client machine, set up an SSH tunnel where 49999 corresponds to the **HPSS_SSM_SERVER_LISTEN_PORT**, the user name is *joe* and the SSM Server machine is "example.com".

```
% ssh -N -f -L 49999:localhost:49999 joe@example.com
```

- On the client machine, run the GUI:

- For Kerberos authentication:

```
% hpssgui.pl -S hpssgui.sessionlog -k krb5.conf -n 49999 -h localhost
```

- For UNIX authentication:

```
% hpssgui.pl -S hpssgui.sessionlog -s unix -u example.com -n 49999 -h
```

localhost

Either the *HPSS Login* window or the SSM main window should open on the client machine for the user. If it doesn't, then retry the last step, running the GUI, using the **-d** option for debug output and the **-S** option to log output to a session log file. This file will provide some information about what is going wrong.

3.4. Multiple SSM sessions

Multiple concurrent sessions of the graphical user interface and command line utility can be executed by the same user with no special modifications. Each session should specify a unique session log file.

3.5. SSM window conventions

This section lists conventions used by SSM and Java, on which SSM is based. The following list does not cover all features of all windows; it only describes the most important points.

- Lists may be sorted by any column. Click on the column header of the desired column to sort the list by the items in that column. The column header will become highlighted and will display an up or down arrow to indicate the direction of the sort. Click the column header a second time to change the direction of the sort.
- List tables have a field that shows the number of displayed and total items in the list in the format "X/Y" where X is the number of items displayed and Y is the total number of items in the list. The field is left-justified under the table. The X and Y values will differ if preferences are set to filter some items out of the list.
- The button panel to the right of the list can be hidden or displayed by clicking the tall, thin button between the list and button panel labeled //. If the button is clicked when the panel is displayed, the button panel will hide, allowing more space for the list. The button panel may be redisplayed by clicking the // button again.
- Colors and fonts are used consistently from window to window. They may differ from platform to platform because the default Java Look and Feel settings vary between platforms.

The **hpssgui** script accepts the following flag parameters in order to control the graphical user interface's look and feel:

-F "Look and Feel"

Valid values: `windows`, `mac`, `motif`, `metal`, and `gtk`. Select the Look and Feel that is applicable to the platform on which the graphical user interface is running. Custom Look and Feels are also available at <https://docs.oracle.com/javase/tutorial/uiswing/lookandfeel/plaf.html>

-b "background color"

The only Look and Feel that supports color settings and themes is the metal Look and Feel. The color may be set by using the color name or hexadecimal Red/Green/Blue value. Here are some examples:

Name	Hexadecimal value
red	0xff0000

green	0x00ff00
blue	0x0000ff
cyan	0x00ffff
yellow	0xffff00
magenta	0xff00ff

-T "theme file"

The theme setting is only applicable when used with the metal Look and Feel. There are eight color parameters that can be used to customize the look of HPSS windows: three primary colors, three secondary colors, black, and white. The color settings may be stored in a theme file using the following syntax:

```
hpss.ssm.ssmuser.hpssgui.CustomMetalTheme.primary1=<color>
hpss.ssm.ssmuser.hpssgui.CustomMetalTheme.primary2=<color>
hpss.ssm.ssmuser.hpssgui.CustomMetalTheme.primary3=<color>
hpss.ssm.ssmuser.hpssgui.CustomMetalTheme.secondary1=<color>
hpss.ssm.ssmuser.hpssgui.CustomMetalTheme.secondary2=<color>
hpss.ssm.ssmuser.hpssgui.CustomMetalTheme.secondary3=<color>
hpss.ssm.ssmuser.hpssgui.CustomMetalTheme.black=<color>
hpss.ssm.ssmuser.hpssgui.CustomMetalTheme.white=<color>
```

<color> should be specified using the color name or Red/Green/Blue hexadecimal value (see the example under the **-b** flag above).

If the theme file location is not specified on the command line, the default value used is `${HOME}/hpss-ssm-prefs/DefaultTheme`.

- Buttons may be "disabled" when the current state of the window does not allow an operation to be performed. In this state, a button is visible but its label text is grayed out and clicking it has no effect. The disabled state occurs when the operation is not supported for the selected item or the SSM user does not have sufficient authority to perform the operation
- A "text" field is any field which displays alphanumeric text or numeric data. This does not include "static" text painted on the window background or labels on things like buttons. Text fields may appear as single or multiple lines and they may be "enterable" (the displayed data can be altered) or "non-enterable" (the displayed data cannot be changed directly).
- Non-enterable text fields have gray backgrounds. A particular field may be enterable under one circumstance but non-enterable under another; for example, a server configuration window's Server ID field is enterable during server creation but may not be changed when modifying a pre-existing configuration record. Additionally, a field is non-enterable when the user does not have sufficient authority to modify the field.
- Enterable text fields have white backgrounds. In most cases, when data in the current field is modified and the field loses focus (the cursor leaves the field), a floppy disk icon will be displayed next to the field to give a visual cue that the field has been changed and that the changes have not been saved. When all changes are made, the user can submit the modifications by clicking one of the window's operation buttons.
- Some enterable text fields are wider than they appear. As typing proceeds and the cursor reaches the right-most edge of the field, the text automatically scrolls to the left and allows further data entry until the actual size of the field has been reached. Scroll back and forth within the field using the left and right cursor keys to view the contents.

- Some text fields which accept integer values can also accept numeric abbreviations such as "KB", "MB", "GB", "TB", or "EB" to specify kilobytes, megabytes, gigabytes, terabytes, or exabytes, respectively. Character case is ignored. For example, entering "1024" will yield the same results as entering "1kb". The entered value must fall within the acceptable numeric ranges for the specified field.
- Some text fields which accept integer values can accept the values in decimal, octal, or hexadecimal form. For these fields, values which begin with an *x* or *0x* will be interpreted as hexadecimal and values which begin with a zero *0* (but not *0x*) will be interpreted as octal. All other values will be interpreted as decimal.
- A combination box is a non-enterable text field inside a button-like box with a small arrow on the right side. Clicking on the box will display a list of items. Selecting a list item will replace the displayed contents of the combination box's text field. Alternately, the list can be dismissed by clicking the mouse anywhere outside of the drop-down list and the displayed contents will remain unchanged.
- A check box is a field containing a box graphic followed by a label. The box may be hollow, indicating that the item is not selected. It may be filled in, or contain a check mark, indicating that the item is selected. Clicking on an enterable check box toggles the state of the selected item. When the state of a check box cannot be modified, it will appear gray in color.
- A radio button is a field containing a circle followed by a label. The circle may be hollow, indicating that the item is not selected or may have a solid interior, indicating that the item is selected. Radio buttons are displayed in groups of two or more items. Only one item within the group can be selected; selecting one button in a group will cause all other buttons in the group to become unselected. When the state of the radio buttons cannot be modified, they will appear gray in color.
- An enterable field containing a cursor is said to have "input focus". If an enterable text field has input focus, typing on the keyboard will enter characters into the field.
- Select, cut, copy, and paste operations can be performed on enterable text fields; on non-enterable fields, only select and copy operations can be performed.
- In some cases, modifying a field value or clicking a button causes the action to be performed immediately. A confirmation window will pop up to inform the user that all changes made to the data window will be processed if the user wishes to continue. If the user selects *No* on the confirmation window, the request will not be processed and any field modifications to the window will continue to be displayed. Some examples are changes to the **Administrative State** field, clicking the Gatekeeper's **Read Site Policy** button, and selecting an entry from the **MPS Storage Class Information Control** combination box.

3.6. Common window elements

Certain SSM buttons and toggle boxes have the same behavior on all SSM windows. Descriptions for these common elements are given below and are not repeated for each window:

Time Created by System Manager field

The last time the System Manager created the structure for this window.

Time Updated by System Manager field

The last time the System Manager updated the data for this window.

Time Received by Client field

The last time the SSM client received an update for this window from the System Manager.

Dismiss button

Closes the current SSM window.

Add button

The **Add** button is displayed on configuration windows when a new configuration record is being created. After the configuration fields are appropriately completed, click the **Add** button to save the data and create the new record. When the Add operation is not permitted, the **Add** button will not be displayed or will appear gray in color.

Update button

The **Update** button is displayed on configuration windows when an existing record is being modified. After the configuration's fields have been modified, click the **Update** button to save the modifications. When the update operation is not permitted, the **Update** button will not be displayed or will appear gray in color.

Delete button

The **Delete** button is displayed on configuration windows of existing records. Click the **Delete** button only when the current record is no longer needed and any dependent records have also been deleted. When the Delete operation is not permitted, the **Delete** button will not be displayed or will appear gray in color.

Start Over button

This resets the current values in a configuration window to the values used when the window was first opened.

Start New button

This replaces the contents of the current configuration window with a new configuration of the same type as the one being viewed. The new configuration's initial values will contain defaults.

Clone (partial) button

This replaces the contents of the current window with a new configuration using some of the current configuration's field values.

Clone (full) button

This replaces the contents of the current window with a new configuration using all of the current configuration's field values.

Freeze

A check box that, while checked, suspends the automatic updates made to an SSM window. This allows reviewing information at the frozen point in time. Unchecking the check box will reactivate normal update behavior.

Refresh button

This requests an immediate update of the displayed information. This can be useful if the user does not wish to wait for an automatic update to occur.

Preferences button and combination box

Edit button

Clicking the **Edit** button opens a configuration window from which certain display characteristics of the parent window can be modified and saved in a preference record. New preference records can be created by saving the preference record with a new name.

Preferences combination box

Click on the Preference combination box to view a list of available preference records used to control the information displayed on the window. The preference record can be modified by either selecting another preference record or by modifying the current preference record. See the **Edit** button description above.

Status Bar

A non-enterable text field along the bottom of all SSM data windows. Status lines display messages concerning the state of the window and the progress of operations started from the window. To view status messages that are longer than the length of the status bar, either stretch the window horizontally or mouse-over the status message to see the full text displayed as a tool tip. Alternately, the user can view status messages in the session log file. When the status bar has had messages written to it, the most recent messages can be viewed in the status bar's tooltip. If there are status messages to view, rolling the mouse over the status bar without clicking gives a tooltip that says, "Click mouse in status bar to view messages". If there are no status messages then the tooltip says, "No status messages". This message stays up for about four seconds or until the user moves the mouse out of the status bar area. To view up to the last 30 messages that have been written to the status bar, click on the status bar. The tooltip that results will show up to the last 30 messages and will remain visible for ten minutes or until the mouse is moved out of the status bar.

File menu

All SSM data windows have the **File** menu. It consists of menu options for controlling the window's location, the user's session, and printing. The **File** menu offers the following options: **Cascade**, **Page Setup**, **Print**, **Close All** or **Close**, **Logoff**, and **Exit**. The **Cascade**, **Close All**, **Logoff**, and **Exit** menu options are only available on the *HPSS Health and Status* window.

Page Setup

SSM uses Java's print facility to create a dialog box enabling the user to enter directions to be sent to the printer. The Page Setup dialog box can be used to specify print media, page orientation, and margin characteristics. The Page Setup dialog box can also be accessed via the Print dialog box (see below). The **Page Setup** menu item is available on all SSM windows.

Print

The Print dialog box is used to set printing options. The print options that are available are platform dependent and therefore may vary. Some print options that can be configured include selecting a printer, setting page size and margins, and specifying the number of copies and pages to print. The **Print** menu item is available on all SSM windows.

Close

The **Close** menu option is used to close the currently selected window. The **Close** menu item is available on all SSM windows.

Edit menu

The **Edit** menu is located on all SSM data windows. From each **Edit** menu, the user can access cut, copy, and paste functions which enable the user to remove data from text fields or transfer data among them. Editable text fields can be updated. Non-editable text fields can be copied, but not changed. Field labels cannot be copied.

Most windowing systems provide keyboard shortcuts for the cut, copy, and paste commands. A typical set of keyboard shortcuts is Ctrl-C for copy, Ctrl-X for cut, and Ctrl-V for paste, but details may vary from system to system. Cut or copied text can be pasted into other applications using the keyboard shortcuts.

- To delete data from a text field: highlight the characters to be removed and select **Cut**.
- To move data from one field to another: highlight the characters to be moved and select **Cut**. Then position the cursor where the data should be placed and select **Paste**.
- To copy data from one field to another: highlight the characters to be copied and select **Copy**. Then position the cursor where the data should be placed and select **Paste**.

Column View menu

The **Column View** menu only appears on SSM windows that display an SSM table. An entry for each column in the table appears in the drop down list along with a corresponding check box. If the check box is selected, then the column will appear in the window's table; otherwise the column will be hidden. Clicking on the check box will toggle the hidden or viewable state of the column.

Help menu

All SSM windows have a **Help** menu. See *Section 3.7, "Help menu overview"* below for detailed information on SSM help.

3.7. Help menu overview

The **Help** menu provides access to online help that is pertinent to the window being displayed. The **Help** menu is available on all SSM data windows but is not available on informational windows such as error messages and confirmation windows. After you select **Help**, the menu will expand to list the help topics that are available for the current window. Selection of a window-related help topic will open an HTML file and jump to the corresponding topic section. Selection of the *HPSS Management Guide* on a help menu will take the user to the table of contents of the management guide.

The *HPSS Management Guide*, along with the *Problem diagnosis and resolution* chapter of the *HPSS Error Manual*, is the main source for diagnosing and solving problems.

In order to access SSM Help, the help files must be installed and accessible to the graphical user interface. For information on obtaining and installing the SSM Help files, see the *Software installation packages* section of the *HPSS Installation Guide* (in the *HPSS installation and infrastructure configuration* chapter).

The SSM Help facility uses two environment variables, `HPSS_HELP_URL_TYPE` and `HPSS_HELP_FILES_PATH`, to determine the location of the SSM Help files. The `HPSS_HELP_URL_TYPE` environment variable specifies the type of URL to aid the browser

in locating the help files. Valid URL types are "https:", "http:", or "file:". The default value for the HPSS_HELP_URL_TYPE is "file:". The HPSS_HELP_FILES_PATH environment variable specifies the location of the installation directory for the SSM Help files. The SSM Help files must exist in HTML format for the graphical user interface to display them. The default value for the HPSS_HELP_FILES_PATH environment variable is "/hpss_src/hpss-<version>/ssmhelp". The values of these environment variables may be overridden.

3.8. Monitor, Operations, and Configure menus overview

The **Monitor**, **Operations**, and **Configure** menus are used by the System Manager to monitor, control, and configure HPSS. They are available only from the *HPSS Health and Status* window. This section provides a brief description for each submenu option listed under the **Monitor**, **Operations**, and **Configure** menus. See related sections for more detailed information about the window that opens after the menu option has been selected.

3.8.1. Monitor menu

The **Monitor** menu contains submenu items that are commonly used by the system administrator to monitor the status of HPSS. The menu is organized with the submenu items that open SSM list windows at the top and the submenu items that open other window types at the bottom.

Alarms & Events

Opens the *Alarms and Events* window which displays HPSS alarm and event messages. Detailed information on each alarm or event can be viewed by selecting an entry from the list and clicking the **Alarm/Event Info** button.

Devices & Drives

Opens the *Devices and Drives* window which displays information on all configured Mover devices and PVL drives. Devices and drives may be created, deleted, viewed, locked, unlocked, dismounted, and marked as repaired from this window.

Tape Drive Quotas

Opens the *Tape Drive Quotas* window which displays the tape drive quota allocations by PVR and tape drive type. The tape drive Recall Limit may be set from this window.

PVL Jobs

Opens the *PVL Job Queue* window which displays all active PVL jobs. Use this window to get more detailed information on a job or to cancel a job.

Servers

Opens the *Servers* window which displays information on all configured HPSS servers. This window can also be used to perform server-related operations such as configuration, startup, shutdown, and viewing server status.

Storage Classes, Active

Opens the *Active Storage Classes* list window which displays information for all storage classes which have been assigned to a Migration/Purge Server and assigned storage resources. Migration,

purge, repack, and reclaim operations can be initiated from this window. This window also provides access to information for each storage class such as the managed object information, configuration record, migration policy, and purge policy. The Migration/Purge Server must be running in order for its storage classes to be active and show up in this list.

RTM Summary

Opens the *RTM Summary List* window which displays summary information for the active Real-Time Monitor (RTM) records. RTM records are maintained by the Core, Gatekeeper, and Mover components of HPSS.

Filesets & Junctions

Opens the *Filesets & Junctions List* window which displays information for the filesets and junctions that are configured in the HPSS system. Filesets and junctions can be created and deleted and details for filesets can be viewed from this window.

Tape Requests

This submenu lists the different tape request list window types.

Check-In

Opens the *Check-In Requests* window which displays all current requests for tape check-ins.

Mount

Opens the *Mount Requests* window which displays all current requests for tape mounts.

Accounting Status

Opens the *Subsystem* list window where the **Accounting Status** and **Start Accounting** buttons can be found.

Lookup HPSS Objects

This submenu lists the type of objects which can be looked up by specifying the object's identifying information.

Cartridges & Volumes

Opens the *Lookup Cartridges and Volumes* window allowing identification of a cartridge or volume by name. A choice can then be made between viewing PVL volume information, PVR cartridge information, or Core Server volume information for the specified cartridge or volume.

Files & Directories

Opens the *Lookup Files and Directories* window into which a pathname can be entered. The pathname can identify a file, directory, or junction. From this window, click the **Show File/Directory** button to display detailed information about the file, directory, or junction.

Objects by SOID

Opens the *Lookup Object by SOID* window where an HPSS object can be specified by entering its HPSS Storage Object ID (HPSS SOID). This window provides access to bitfile and virtual volume information.

SSM Information

This submenu lists the options available for viewing statistics for the System Manager and the user client session.

System Manager Statistics

Opens the *SSM System Manager Statistics* window to view statistics for the System Manager such as the number of RPC calls, notifications, and messages that were processed. **User Session Information** opens the *User Session Information* window to display the user's login name, authority, and statistics regarding the user's session.

3.8.2. Operations menu

Accounting Report

Opens the *Subsystems* list window where a subsystem can be highlighted and the **Start Accounting** button can be selected to obtain an accounting report.

Drive Dismount

Opens the *Devices and Drives* list window where the **Dismount Drive** button is located.

PVL Job Cancellation

Opens the *PVL Job Queue* window from which PVL jobs may be selected and canceled.

Resources

This submenu lists the operations that can be performed on disk and tape cartridges and volumes.

Import Disk Volumes

Opens the *Import Disk Volumes* window where a list of disk volume labels can be entered and an import request can be submitted.

Import Tape Volumes

Opens the *Import Tape Volumes* window where a list of tape volume labels can be entered and an import request can be submitted.

Create Disk Resources

Opens the *Create Disk Resources* window where a list of disk volume labels can be entered and a request to add the disks to a storage class can be submitted.

Create Tape Resources

Opens the *Create Tape Resources* window where a list of tape volume labels can be entered and a request to add the tapes to a storage class can be submitted.

Delete Resources

Opens the *Delete Resources* window allowing deletion of existing tape or disk storage resources.

Export Volumes

Opens the *Export Volumes* window which exports tape cartridges and disk volumes, making them unavailable to HPSS.

Move Cartridges

Opens the *Move Cartridges To New PVR* window allowing ownership of tape cartridges to be transferred between PVRs.

Migrate/Purge Data

Opens the *Active Storage Classes* window. From this window, a storage class may be highlighted and a migration or purge can be started by selecting the corresponding button.

Repack/Reclaim Tapes

Opens the *Active Storage Classes* window where a storage class may be highlighted and the **Repack Volumes** or **Reclaim Volumes** button selected to perform the operation.

Ping System Manager

Selecting this submenu option tests the connectivity between the GUI and the System Manager. If the ping is successful, nothing will happen. If the ping is unsuccessful, an error message will be displayed.

Shutdown

This submenu provides a quick way to send a shutdown request to any server other than the Startup Daemon. If you want to shut down a particular server or set of servers, use the **Shutdown** or **Force Halt** buttons on the *Servers* list window. The System Manager cannot be shutdown via the *Servers* list window. The Startup Daemon cannot be shutdown at all using SSM.

All Non-SSM Servers

Sends a shutdown command to all servers other than the System Manager and Startup Daemon. Note that some servers may take a few minutes to shut down. To restart the servers, select the servers in the *Servers* list window and click the **Start** button.

System Manager

Sends a shutdown command to only the System Manager.

3.8.3. Configure menu



*It is recommended that the system administrator configure an HPSS system traversing the **Configure** menu in top-down order since some configuration items have a dependency on others.*

Subsystems

This opens the *Subsystems* list window where a list of all configured subsystems can be viewed, new subsystems can be configured, or existing subsystems can be deleted. Additionally, from the *Subsystems* list window, accounting statistics can be viewed and reports can be generated.

Policies

This submenu lists the policy types that can be configured for HPSS.

Accounting

Opens the *Accounting Policy* window allowing configuration and management of the accounting policy. Only one accounting policy is allowed.

Location

Opens the *Location Policy* window allowing configuration and management of the location policy. Only one location policy is allowed.

Logging

Opens the *Logging Policies* list window allowing configuration and management of the logging policies.

Migration

Opens the *Migration Policies* list window allowing configuration and management of the migration policies.

Purge

Opens the *Purge Policies* list window allowing configuration and management of the purge policies.

Storage Space

This submenu lists the storage space configurations required for HPSS. Classes of Service contain a hierarchy, and hierarchies are made up of a list of storage classes.

Storage Classes

Opens the *Configured Storage Classes* list window allowing configuration and management of storage classes.

Hierarchies

Opens the *Hierarchies* list window allowing configuration and management of storage hierarchies.

Classes of Service

Opens the *Class of Service* list window allowing configuration and management of the Classes of Service.

Servers

This opens the *Servers* list window, which will facilitate server configuration and management.

Global

This opens the *Global Configuration* window allowing the configuration and management of the HPSS global configuration record. Only one global configuration is allowed.

Devices & Drives

This opens the *Devices and Drives* list window allowing configuration and management of devices and drives for use with HPSS.

File Families

This opens the *File Families* list window allowing the configuration and management of file families.

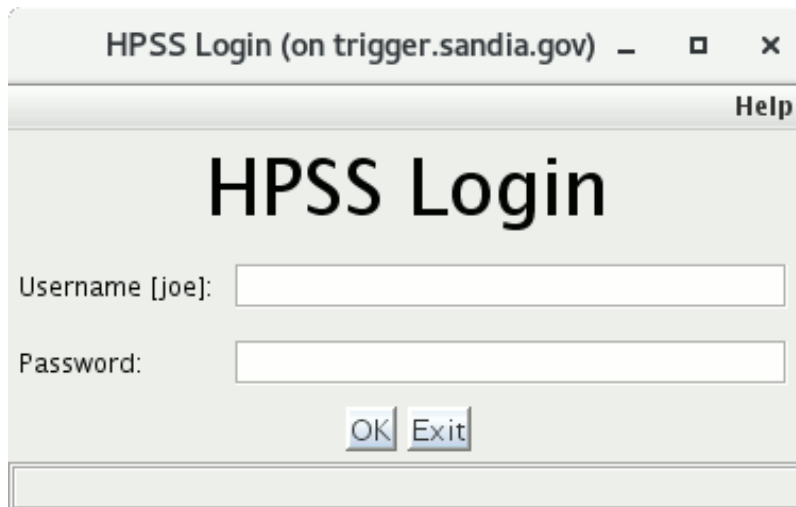
List Preferences

This submenu contains an entry for each SSM list window. A preference record can be created or the default preference record can be modified to allow the user to customize each SSM list window's data view by using filtering methods. See *Section 3.10, "SSM list preferences"* for more information.

3.9. SSM-specific windows

This section describes the *HPSS Login* window, the *HPSS Health and Status* window, and the SSM information windows.

3.9.1. HPSS Login



The *HPSS Login* window may appear after starting the **hpssgui** script. The user must supply a valid HPSS user name and password in order to access SSM and monitor HPSS. The SSM GUI will attempt to use existing credentials (for example, Kerberos with kinit) first, a keytab next, and falling back to a username/password. If the login window appears, it means that the SSM GUI did not find valid credentials from the first two methods.

If a login attempt is unsuccessful, review the user session log for an indication of the problem. See the **hpssadm** or **hpssgui** man pages for more information about the user session log.

Field descriptions.

User ID

If this field appears, it will be in a form dictated by the `login.conf` mechanism. Enter a valid user ID here.

Password

Some variation of this prompt will appear, as dictated by the `login.conf` mechanism. Enter the password for the user ID.

Note that one or both of these prompts may appear. For instance, the username may be requested first, followed by another window asking for the password. The particulars are up to the mechanism specified in `login.conf`.

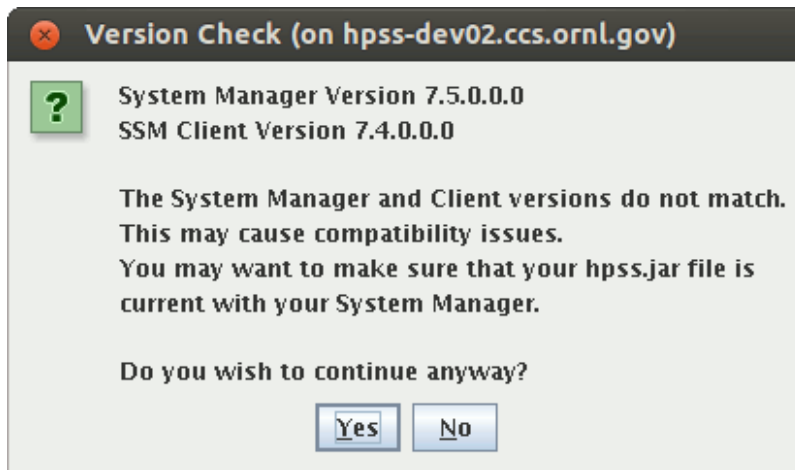
OK

Attempt to contact the System Manager and authenticate the user. If the attempt is successful, the *HPSS Health and Status* window will open. Otherwise, an error message will be displayed in the status bar.

Exit

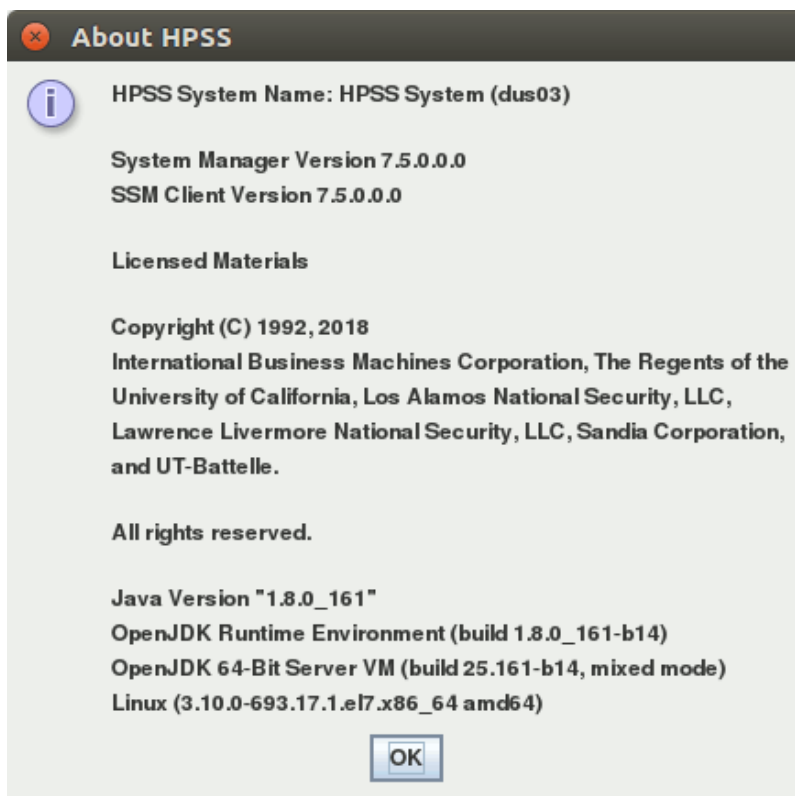
Close the *HPSS Login* window and terminate the **hpssgui** client session.

If the System Manager and SSM Client versions do not match then the following window will be displayed after the user logs in:



The user may choose to continue logging in or to exit. However, as the dialog says, running with mismatched versions may cause compatibility problems.

3.9.2. About HPSS

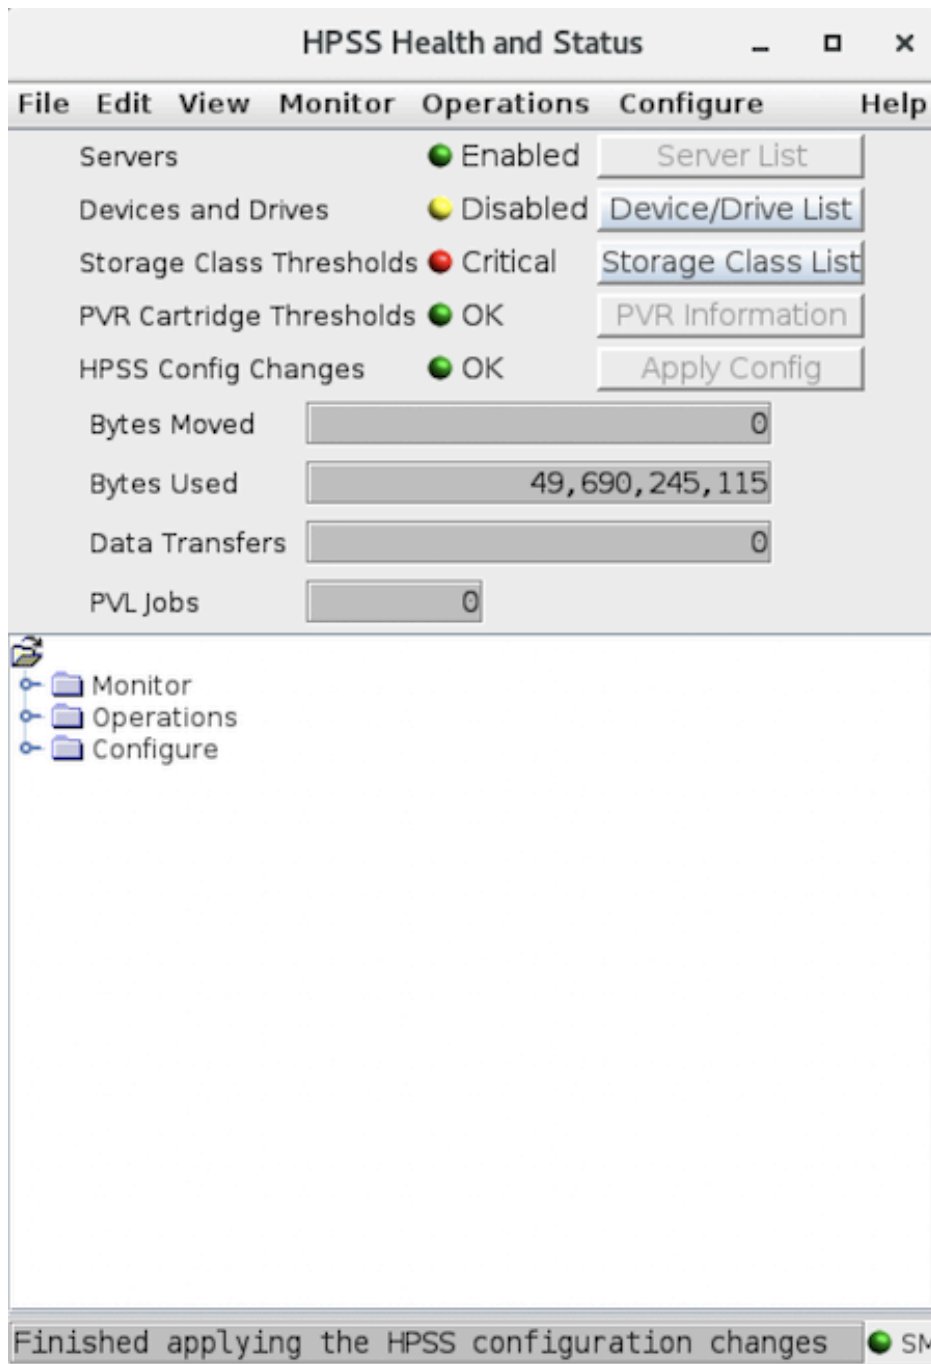


The *About HPSS* window displays version information and a portion of the HPSS copyright statement. The *About HPSS* window is accessible by selecting the Help menu's "About HPSS" submenu from any of the **hpssgui** windows.

The HPSS System Name and System Manager Version are not displayed when the *About HPSS* window is requested from the *HPSS Login* window. These two pieces of information are not available until the user actually logs into the System Manager.

Differences in the System Manager and SSM Client versions may indicate that the code for the client or System Manager (or both) should be updated.

3.9.3. HPSS Health and Status



When a user successfully connects to the System Manager through the *HPSS Login* window, the *HPSS Health and Status* window replaces the *HPSS Login* window on the screen. The *HPSS Health and Status* window will remain on the screen until the user exits or logs out. It provides the main menu and displays information about the overall status of HPSS.

The *HPSS Health and Status* window is composed of several high-level components, each of which is discussed in its own section below.

3.9.3.1. SM server connection status indicator

The SM connection status indicator is located in the bottom right corner of the status bar. When the SM icon is red, the client's connection to the System Manager is lost; when it is green, the connection is active.

3.9.3.2. HPSS status

On the upper section of the *HPSS Health and Status* window are five status fields that represent the aggregate status of the HPSS system:

Servers

This displays the most severe status reported to SSM by any HPSS server.

Devices and Drives

This displays the most severe status as reported to SSM for any configured Mover device or PVL drive.

Storage Class Thresholds

This displays the most severe status reported by the MPS for the Active Storage Class space usage. SSM assumes that all thresholds are "OK" until it receives contrary information.

PVR Cartridge Thresholds

This displays the most severe status of cartridge usage reported by any configured PVR. SSM assumes that all thresholds are "OK" until it receives contrary information.

HPSS Config Changes

This displays whether certain configuration changes have been made that require server(s) to be reinitialized. Clicking the **Apply Config** button when it is enabled will reinitialize the server(s) to apply the configuration changes.

For the **Servers** and **Devices and Drives** fields, possible status values in increasing order of severity are:

- **Normal** indicates no problem reported.
- **Unknown** indicates SSM cannot determine the true status due to communication problems or other difficulties.
- **Suspect** indicates a situation has been encountered by a server or device which might be a problem, but does not significantly affect HPSS operation.
- **Minor** indicates a problem has been encountered by the server or device, but it does not significantly affect HPSS operation.
- **Major** indicates a problem has been encountered by the server or device that may degrade HPSS operation.
- **Critical** indicates a problem has been encountered by the server or device that may disrupt HPSS operation until the problem is resolved.

For the **Storage Class Thresholds** field, the possible status values are:

- `OK` indicates no problem reported.
- `Warning` indicates a threshold has been reported as crossing its warning limit.
- `Critical` indicates a threshold has been reported as crossing its critical limit.
- `Stale` indicates the Migration/Purge Server is down or SSM has not received an update from the MPS.

For the **PVR Cartridge Thresholds** field, the possible status values are:

- `OK` indicates no problem reported.
- `Warning` indicates a threshold has been reported as crossing its warning limit.
- `Critical` indicates a threshold has been reported as crossing its critical limit.
- `Unknown` or `Stale` indicates the PVR is down or SSM has not received an update from the PVR.

For the **HPSS Config Changes** field, the possible status values are:

- `OK` indicates no configuration changes need to be applied.
- `Pending` indicates pending configuration changes need to be applied.
- `Error` indicates an error occurred applying the configuration changes.
- `Unknown` indicates the SSM cannot determine if the configuration changes were applied due to a communication problem or other difficulties.

As problems are resolved or returned to normal, the status fields will automatically reflect the changes.

In addition to the text which describes the status, these fields are displayed with colored icons. The icon color depicts that status as follows:

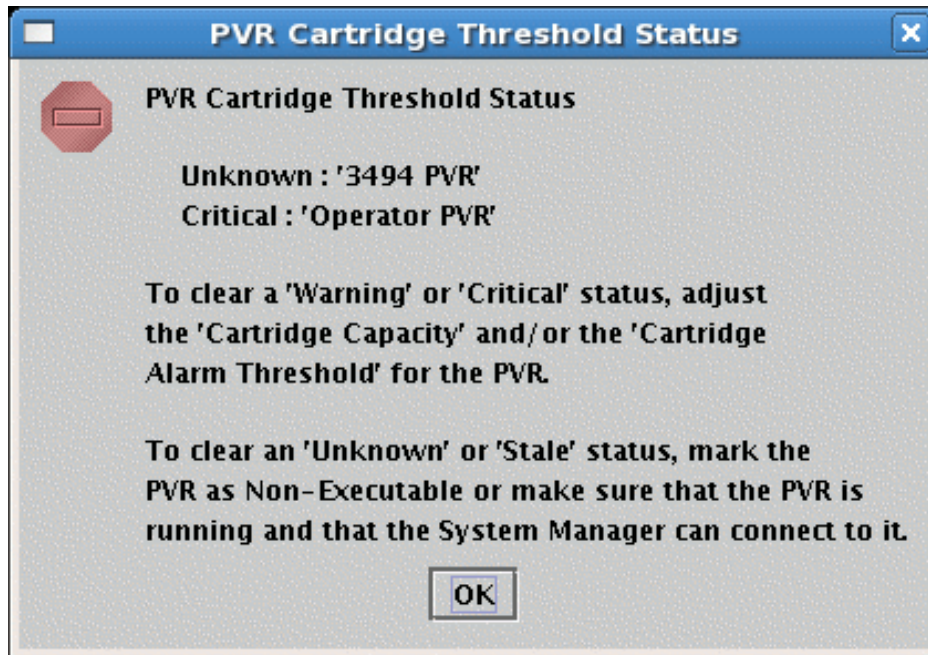
- Red for major and critical problems
- Magenta for minor problems
- Yellow for unknown, stale, suspect, and warning problems
- Green for normal; no problem

Click on the button to the right of the status icon to get more details.

For **Servers, Devices and Drives**, and **Storage Class Thresholds** the button will open the corresponding SSM list window in sick list mode. Once this window is open, use it to get detailed status information on the sick entries, assuming the HPSS servers are still healthy enough to respond to SSM requests.

See *Section 3.10, “SSM list preferences”* for more information on the sick list mode.

For **PVR Cartridge Thresholds** the button will display a message dialog with information about the PVRs that have cartridge threshold issues. This message dialog will look like the following:



The status section of the *HPSS Health and Status* window can be hidden from view by selecting the **View** menu item and unchecking the **HPSS Status** check box.

3.9.3.3. HPSS statistics

The HPSS statistics fields are located in the middle section of the *HPSS Health and Status* window and display the number of bytes moved, bytes used, data transfers, and current PVL jobs in the system. The number of bytes moved and number of data transfers indicate the data accumulated by the Movers since startup or data reset.

HPSS statistics fields show general trends in HPSS operations; the numbers are not all-inclusive. Some values may fluctuate up and down as servers are started or shut down. Some values, such as **Bytes Moved**, can be reset to zero in individual Movers and by SSM users.

Bytes Moved

Total bytes moved as reported by all running Movers.

Bytes Used

Total bytes stored on all disk and tape volumes as reported by all running Core Servers.

Data Transfers

Total data transfers as reported by all running Movers.

PVL Jobs

Total jobs reported by the PVL.

The statistics section of the *HPSS Health and Status* window can be hidden from view by selecting the **View** menu item and unchecking the **HPSS Statistics** check box.

3.9.3.4. Menu tree

The menu tree section of the *HPSS Health and Status* window displays a tree structure that mirrors the structure of the **Monitor**, **Operations**, and **Configure** menus. The menu tree can be fully expanded so

that the user can view all the menu options for monitoring, configuring, or operating HPSS. Selecting a leaf of the menu tree results in the same response as selecting the same item from the *HPSS Health and Status* menu bar. The user can also choose to expand only those branches of the menu tree that are of interest. The collapsed and expanded state of the branches on the menu tree can be toggled by clicking on the branch icons.

The menu tree can be hidden from view by selecting the **View** menu item and unchecking the **Menu Tree** check box.

3.9.3.5. File menu

All SSM data windows have the **File** menu. See *Section 3.6, “Common window elements”* for a description of the **File** menu options that appear on all SSM data windows. The following **File** menu options are unique to the *HPSS Health and Status* window:

Cascade

Select the **Cascade** menu option to rearrange the SSM windows, placing them on the screen starting in the upper left-hand corner and cascading downward toward the lower right-hand corner of the user’s desktop. When the cascading has been completed, the *HPSS Health and Status* window will be centered on the screen and brought to the foreground.

Close All

Select **Close All** to close all SSM windows except the *HPSS Health and Status* window. To close the *HPSS Health and Status* window, the user must select **Logoff** or **Exit**.

Logoff

Select **Logoff** to exit the SSM session and close all SSM windows. The **hpssgui** script will continue to be active and will still use the same user’s session log. This will begin the login process again. Note, however, that if the login process used a keytab or existing credentials, this will automatically log the user back in. If this was the case, it is better to choose the **Exit** option instead.

Exit

Select **Exit** to exit the SSM session and close all SSM windows. The **hpssgui** script will terminate and the user’s session log will be closed. The user must rerun the **hpssgui** script to access the *HPSS Login* window and reconnect to the System Manager.

3.9.3.6. View menu

The **View** menu is only available on the *HPSS Health and Status* window. The **View** menu offers the user the ability to hide or display elements of the *HPSS Health and Status* window in order to optimize the viewable area. Under the **View** menu there is a menu item and check box for each window element that can be hidden. If the box contains a check mark then the corresponding section of the *HPSS Health and Status* window that displays this element will be visible. If the check box is empty, then the element is hidden from the window view. Clicking on the check box will toggle the visible state of the window element.

HPSS Status

Toggle the **HPSS Status** menu option to hide or view the HPSS status section of the *HPSS Health and Status* window.

HPSS Statistics

Toggle the **HPSS Statistics** menu option to hide or view the HPSS statistics section of the *HPSS Health and Status* window.

Menu Tree

Toggle the **Menu Tree** menu option to hide or view the menu tree section of the *HPSS Health and Status* window.

3.9.4. SSM information windows

These windows describe the System Manager and the user's **hpssgui** session.

3.9.4.1. System Manager Statistics window

This window displays the statistics about the number of RPC calls, notifications, and messages that the System Manager has processed since its Start Time. While the window is open, the statistics are updated at timed intervals. The information in this window is intended for use in analyzing System Manager performance or in troubleshooting. It is of limited use in everyday production situations.

To open the window, start at the *HPSS Health and Status* window. From its menu bar, select **Monitor > SSM Information > System Manager Statistics**.

SSM System Manager Statistics

File Edit Column View Help

Start Time: Apr 16, 2018 5:30:45 PM

Uptime: 00-20:45:53

CPU Time: 00-00:00:21

Memory Usage: 3559120

Process ID: 17937

Hostname: hpss-dev-dus03.ccs.ornl.gov

RPC Calls to Servers: 1788

Alarm Messages: 127

Event Messages: 11764

RPC Interface Information

	Server	Client
Status	OK	OK
Thread Pool Size	100	100
Request Queue Size	20	600
Active RPCs	0	18
Queued RPCs	0	0
Maximum Active/Queued RPCs	4	20
Data Change Notifications	15960	
Unsolicited Notifications	0	
Tape Check-In Messages	0	
Tape Mount Messages	266	
Total RPCs	16226	53010

Client Connection Information

Maximum Connections: 4

Current Connection Count: 4

Current Client Count: 2

ID	State	User Auth	Principal	Hostname
0	Free	admin	dusadmin@hpss-dev-dus03.ccs.ornl.gov	hpss-dev-dus03.ccs.ornl.gov
1	Free	admin	dusadmin@hpss-dev-dus03.ccs.ornl.gov	hpss-dev-dus03.ccs.ornl.gov
2	Free	admin	dusadmin@hpss-dev-dus03.ccs.ornl.gov	hpss-dev-dus03.ccs.ornl.gov
3	Free	admin	dusadmin@hpss-dev-dus03.ccs.ornl.gov	hpss-dev-dus03.ccs.ornl.gov
4	Free	admin	dusadmin@hpss-dev-dus03.ccs.ornl.gov	hpss-dev-dus03.ccs.ornl.gov
5	Free	admin	dusadmin@hpss-dev-dus03.ccs.ornl.gov	hpss-dev-dus03.ccs.ornl.gov
6	Free	admin	dusadmin@hpss-dev-dus03.ccs.ornl.gov	hpss-dev-dus03.ccs.ornl.gov

16/16

☐ Freeze Refresh Dismiss

Preferences

Edit

Default

Field descriptions.

Start Time

The time the System Manager was started.

Uptime

The elapsed clock time since the System Manager was started.

CPU Time

The amount of CPU time that the System Manager has consumed.

Memory Usage

The amount of memory that the System Manager is currently occupying, in kilobytes (KB).

Process ID

The process ID of the System Manager.

Hostname

The name of the host where the System Manager is running.

RPC Calls to Servers

The number of RPCs the System Manager has made to other HPSS servers.

Alarm Messages

The number of alarm-type messages read from the HPSS syslog file.

Event Messages

The number of event-type messages read from the HPSS syslog file.

RPC Interface Information

Provides information about the server and client RPC interfaces. The server interface is used by other HPSS servers to contact the System Manager. The client interface is used by the **hpssgui** and **hpssadm** programs to contact the System Manager. There are two columns of data, one for the server interface and one for the client interface. Not all fields are available for both interfaces. The fields include:

Status

Current status of the thread pool and request queue of the RPC interface. The Status can be:

- **OK** indicates the number of active RPCs is less than the **Thread Pool Size**. There are enough threads in the thread pool to handle all current RPCs with spare ones left over.
- **Warning** indicates the number of active RPCs is greater than the **Thread Pool Size**. The number of queued RPCs is less than 90% of the **Request Queue Size**. There aren't enough threads to handle all the current RPCs and some are having to wait in the queue, but the queue is big enough to hold all the waiters.
- **Critical** indicates the number of queued RPCs is greater than or equal to 90% of the **Request Queue Size**. There aren't enough threads to handle all the current RPCs, some are having to wait in the queue, and the queue is getting dangerously close to overflowing, at which point any new RPCs will be rejected.

Thread Pool Size

The maximum number of RPCs that can be active at any one time. For the server RPC interface this value is determined by the `HPSS_SM_SRV_TPOOL_SIZE` environment variable. For the client RPC interface this value is determined by the **Thread Pool Size** field defined on the *Core Server Configuration* window. Refer to *Section 5.2.1.3, "Interface controls"*.

Request Queue Size

The maximum number of RPC requests that can be queued and waiting to become active. For the server RPC interface this value is determined by the HPSS_SM_SRV_QUEUE_SIZE environment variable. For the client RPC interface this value is determined by the **Request Queue Size** field on the *Core Server Configuration* window. Refer to *Section 5.2.1.3, "Interface controls"*.

Active RPCs

The number of RPCs that are currently active. To be active an RPC must have been assigned to a thread in the thread pool.

Queued RPCs

The number of RPCs that are waiting in the request queue to be assigned to a thread in the thread pool.

Maximum Active/Queued RPCs

The maximum number of RPC requests that were active (in the thread pool) or queued (in the request queue) at the same time. This value can be used to help tune the **Thread Pool Size** and **Request Queue Size** for the RPC interface. If the **Maximum Active/Queued RPCs** is greater than the **Thread Pool Size** you might consider increasing the **Thread Pool Size**, the **Request Queue Size**, or both, to help with the System Manager performance. However, increasing these two parameters could cause the System Manager to require more memory.

Data Change Notifications

The number of data change notifications received from servers. (Server RPC Interface only).

Unsolicited Notifications

The number of notifications which the System Manager received from other HPSS servers but which it did not request from them. (Server RPC Interface only).

Tape Check-In Messages

The number of tape check-in request notifications received. (Server RPC Interface only).

Tape Mount Messages

The number of tape mount request notifications received. (Server RPC Interface only).

Total RPCs

Total number of RPCs processed by the RPC interface.

Client Connection Information

Provides information about clients that have connected to the System Manager.

Maximum Connections

The maximum number of client connections that the System Manager was handling at any one time. Each client can have multiple connections. The default connections per client is "2". Each client can specify the number of connections using the -Dhpss.ssm.SMConnections Java command line option.

Current Connection Count

The number of client connections currently being handled by the System Manager.

Current Client Count

The number of clients currently connected to the System Manager. This will be the number of entries in the **Client List** (below) with an `In Use` state.

Client List

The list of clients connected to the System Manager. The entries that appear in the client list will be for (up to) the last 64 (`SSM_CLIENT_MAX`) clients that have connected to the System Manager.

ID

Slot identifier.

State

Clients displayed in the client list can have one of two states: `In Use` and `Free`.

- `In Use` indicates the clients that are currently connected to the System Manager.
- `Free` indicates the clients that were once connected but are no longer connected (active).

Once the client list contains `SSM_CLIENT_MAX` `In Use` and/or `Free` entries, then the oldest `Free` slot (the client that has been disconnected the longest) is given to the next client. Once all the slots are `In Use`, then the client table is full; no new clients can connect until one of the `In Use` slots becomes `Free` after a client disconnects.

User Auth

The client user's authorization level: `admin` or `operator`. See *Section 3.3, "Configuration and startup of **hpssgui** and **hpssadm**"* for more details.

Principal

The user's principal name.

Hostname

The name of the host where the client is running.

Connections

The number of RPC connections this client has to the System Manager.

Start Time

The time that the client connected to the System Manager.

Connect Time

The elapsed time since the client connected to the System Manager.

Idle Time

The elapsed time since the System Manager received an RPC from the client.

Cred Refreshes

The number of times the principal's credentials have been refreshed since the client has been connected to the System Manager.

RPC Calls

The number of RPCs that the client has made to the System Manager.

RPC Waits

The number of RPCs that the client has currently "waiting" in the System Manager. These are the RPCs which are active and connected to the System Manager but which have not yet completed.

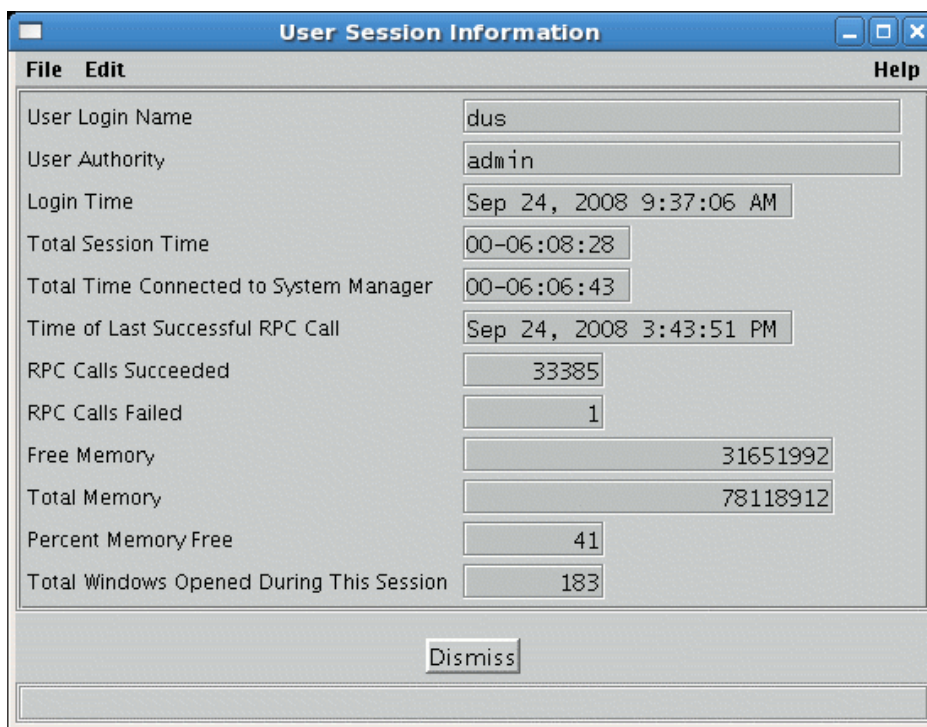
Client UUID

The UUID for the client.

3.9.4.2. User Session Information window

The *User Session Information* window displays memory and network statistics for the current SSM user session. The window is updated in timed intervals as data changes.

To open the window, start at the *HPSS Health and Status* window. From its menu, select **Monitor > SSM Information > User Session Information**.



The screenshot shows a window titled "User Session Information" with a menu bar containing "File", "Edit", and "Help". The window displays various session statistics in a table-like format with labels on the left and values in text boxes on the right. At the bottom, there is a "Dismiss" button.

Field	Value
User Login Name	dus
User Authority	admin
Login Time	Sep 24, 2008 9:37:06 AM
Total Session Time	00-06:08:28
Total Time Connected to System Manager	00-06:06:43
Time of Last Successful RPC Call	Sep 24, 2008 3:43:51 PM
RPC Calls Succeeded	33385
RPC Calls Failed	1
Free Memory	31651992
Total Memory	78118912
Percent Memory Free	41
Total Windows Opened During This Session	183

Field descriptions.**User Login Name**

The name that the user entered as the user ID when logging into HPSS.

User Authority

The authority level that the user has in order to perform operations in SSM. This can be admin or operator.

Login Time

The time that the user logged into SSM.

Total Session Time

The amount of time that the user has had the graphical user interface running.

Total Time Connected to System Manager

The amount of time that the user has been connected to the System Manager. This is the time that has elapsed since the login was successful.

Time of Last Successful RPC Call

The time that the last RPC call was made to the System Manager which completed successfully.

RPC Calls Succeeded

The number of successful RPC calls made by this user session.

RPC Calls Failed

The number of unsuccessful RPC calls made by this user session.

Free Memory

Amount of free memory available in the **hpssgui** program.

Total Memory

Amount of total memory available in the **hpssgui** program.

Percent Memory Free

The ratio of free memory to total memory in the **hpssgui** process.

Total Windows Opened During This Session

The number of windows created during the current user session.

3.10. SSM list preferences

When a user logs into SSM, the System Manager reads the saved preferences file and loads them into the SSM session, if they exist.

Each SSM list type has a Default preferences record. The Default preferences configuration is set so that the more commonly used columns are displayed. The following lists have a preferences record:

- Alarms & Events
- Classes of Service
- Devices & Drives
- File Families
- Filesets & Junctions
- Hierarchies
- Logging Policies
- Migration Policies
- Purge Policies
- PVL Jobs

- RTM Summary
- Servers
- SM Clients
- Storage Classes, Active
- Storage Classes, Configured
- Subsystems
- Tape Check-Ins
- Tape Drive Quotas
- Tape Mounts

The *Servers*, *Devices and Drives*, and *Active Storage Classes* list windows also have Sick preferences which display all items with abnormal status.

Default preferences can be modified but cannot be deleted. Sick preferences may not be modified or deleted.

Preferences are saved on the client node in the directory specified by the "configuration path" **hpssgui** command line argument ("-i"). This option can also be set using the environment variable `HPSSGUI_USER_CFG_PATH` or the configuration file entry `HPSS_SSM_USER_PREF_PATH`. If this option is not specified, the default value is `<client node>:<user.home>/hpss-ssm-prefs`. The user must have permissions to create the preferences file in the directory.

Preferences windows contain filters for controlling the data displayed in the list window. Columns in the list window can be rearranged and resized by dragging columns and column borders on the list window itself. Data in the list can be sorted ascending or descending based on any column in the list by clicking in that column's header. Such changes are kept up to date in the preferences file, without the need to click on a **Save** button.

Columns can be hidden or redisplayed through the **Column View** menu on each list window.

The *Tape Mount Requests* and *Tape Check-In Requests* windows have **Auto Popup** check boxes. The states of these check boxes, as well as those in the **View** Menu of the *HPSS Health and Status* window are stored in the preferences file.

If a user's preferences file becomes obsolete, the current version of the user interface software will convert the old preferences to the current format.

Check box filters.

Check box filters apply to columns that have a limited set of display values. The check box filters are grouped by column name and contain a check box for each allowed value. If the check box is selected, all rows containing the value will be displayed on the corresponding SSM list window. At least one check box filter must be selected in each group.

Text filters.

Columns that can be filtered by their text content are listed in the Text Filters section of the preference window. Users can control which rows are displayed by entering a Java regular expression into one or more of the **Text Filter** fields. If the **Text Filter** field is blank then no filtering on the field will occur.

The preference window will verify that the text entered is a valid Java regular expression before allowing the user to save the preferences.

For example, if the list contains storage class names as follows

```
SClass_4-way_Disk  
SClass_4-way_Tape  
Sclass_4-way_Archive  
Sclass_Disk
```

entering a Java regular expression of **".4-way."** would result in the "Sclass_Disk" entry to be filtered out of the display.

For a complete description on how to use regular expressions in Java, Refer to the following web page: <http://java.sun.com/j2se/1.5.0/docs/api/java/util/regex/Pattern.html>.

Button options.

Save as

Save the current preference settings to a new named preference. A pop-up window will be displayed so that the new name can be entered.

Save

Save the current preference settings to the preferences file using the same preference name.

Delete

Delete the currently displayed preference settings from the preferences file. Sick and Default preference settings cannot be deleted.

Reload

Reread the preference settings from the preferences file.

Apply

Apply the current preference configuration to the parent SSM list window. Clicking **Apply** does not save the changes to the preferences file. To save the preference settings, click the **Save** or **Save as** button. The **Apply** button is only available when the Preferences window is accessed via its parent SSM List window (that is, it is not available when the Preferences window is accessed from the **Configure** menu).

Show List

Brings the window to which these preferences apply to the foreground. The **Show List** button is only available when the Preferences window is accessed via its parent SSM List window.

Dismiss

Close the window.

To modify an existing preference configuration, select the Preference Name from the drop down list and click the **Reload** button. Modify the configuration then click the **Save** button. To create a new

Preference configuration, open an existing configuration, make the modifications, click the **Save as** button then enter a unique value for the new preference name and click the **OK** button.

Chapter 4. Global and subsystem configuration

This chapter discusses two levels of system configuration: global and storage subsystem. The global configuration applies to the entire HPSS installation while subsystem configurations apply only to servers and resources allocated to storage subsystems.

For new HPSS systems, it is recommended that the first step of configuration be the partial definition of subsystems. The subsystems should be defined at this point except for their Gatekeeper and default COS information. The global configuration should be created after server configuration is complete. Last of all, the subsystem configuration is completed by adding the definitions of the Gatekeeper and default COS. See *Section 1.3, “HPSS configuration roadmap (new HPSS sites)”*.

4.1. Global Configuration window

The screenshot shows a window titled "Global Configuration" with a menu bar containing "File", "Edit", and "Help". The window contains several configuration fields and sections:

- System Name:** A text field containing "HPSS System (dus03)".
- HPSS ID:** A text field containing "1".
- Hash Algorithm Seed:** A text field containing "763460996".
- Root Core Server:** A dropdown menu showing "Core Server (Subsystem 1)".
- Default Class of Service:** A dropdown menu showing "1 | Class of Service 1".
- Metadata Space Warning Threshold:** A text field containing "85" followed by "percent".
- Metadata Space Critical Threshold:** A text field containing "90" followed by "percent".
- Metadata Space Monitor Interval:** A text field containing "1800" followed by "seconds".
- DB Log Monitor Interval:** A text field containing "300" followed by "seconds".
- Root User ID:** A text field containing "0".
- COS Change Stream Count:** A text field containing "1".
- Global Flags:** A section containing four checkboxes:
 - ☒ Root is Superuser
 - ☐ Can change UID to self if has Control Perm
 - ☐ Can change UID if has Delete Perm on Security ACL
 - ☐ Object names can contain unprintable characters
- Trashcan Settings:** A section containing:
 - ☐ Trashcans Enabled
 - Number of Trashcan Threads:** A text field containing "5".
 - Trash Incinerator Interval:** A text field containing "10" followed by "seconds".
 - Trash Unlink Eligible Time:** A text field containing "10" followed by "seconds".
 - Trash Statistics Interval:** A text field containing "60" followed by "seconds".

At the bottom of the window, there are three buttons: "Update", "Start Over", and "Dismiss".

This window allows you to manage the HPSS global configuration record. Only one such record is permitted per HPSS installation. To open the window, on the *HPSS Health and Status* window select the **Configure** menu, and from there the **Global** menu item.

Field descriptions.

System Name

An ASCII text string representing the name for this HPSS system.

HPSS ID

The unique ID assigned to each HPSS installation.

Hash Algorithm Seed

The seed value used to define the system's partition hashing algorithm. The partition hashing algorithm is used to determine the database partition in which the metadata for each object (such as, name space object, bitfile, and storage segment) is stored.



Once defined, the **Hash Algorithm Seed** must never be changed.

Root Core Server

The name of the Core Server which manages the root fileset ("/") in the HPSS name space.



Advice

The root Core Server should be selected with care. Once it is chosen it cannot be changed as long as the chosen server exists. If the root Core Server must be changed, the current root Core Server will have to be deleted before SSM will allow another root Core Server to be selected.



The following settings are global defaults. If they are changed, the Core Servers which do not override these defaults will all need to be reinitialized or restarted in order to actualize the new settings. It is recommended to use the **Apply Config** button on the *HPSS Health and Status* window to reinitialize the server(s).

Default Class of Service

The name of the default Class of Service (COS). Core Servers will store new files in this COS when the HPSS client does not specify a COS or provide hints on the creation request. This default can be overridden by each storage subsystem.



Advice

If the COS chosen for the default Class of Service is deleted, be sure to change the default Class of Service before deleting the COS.

If disk is used as the default COS, it should be defined so that a large range of file sizes can be handled. For this reason it may be reasonable to set the default COS to a tape COS.

Metadata Space Warning Threshold

Provides a default value for the metadata warning threshold. When the space used in any DB2 tablespace exceeds this percentage, the Core Server will issue a periodic warning message. This value can be overridden for each storage class on a storage subsystem basis (see *Section 4.2.1, "Subsystems list window"* for more information).

Metadata Space Critical Threshold

Provides a default value for the metadata critical threshold. When the space used in any DB2 tablespace exceeds this percentage, the Core Server will issue a periodic critical alarm. This value can be overridden for each storage class on a storage subsystem basis (see *Section 4.2.1, "Subsystems list window"* for more information).

Metadata Space Monitor Interval

The Core Server will check metadata usage statistics at the indicated interval, specified in seconds. The minimum value for this field is 1800 seconds (30 minutes). A value of zero will turn off metadata space monitoring. This field may be overridden on the Storage Subsystem configuration.

DB Log Monitor Interval

The Core Server will check the consistency of Database Logs and Backup Logs at the indicated interval, specified in seconds. The logs are consistent if both primary and backup log directories exist and contain log files with the same names. The minimum value for this field is 300 seconds (five minutes). A value of zero will turn off DB Log monitoring. This field may be overridden on the Storage Subsystem configuration.

Root User ID

The UID of the user who has root access privileges to the HPSS name space. This only applies if the **Root Is Superuser** flag is set.

COS Change Stream Count

The number of background threads that run in the Core Server to process Class of Service change requests. This field may be overridden on the Storage Subsystem configuration.

Global flags

Root Is Superuser

If checked, root privileges are enabled for the UID specified in the **Root User ID** field. If the box is not checked, the UID specified in the **Root User ID** field will not have root privileges. Root access privileges grant the specified user the same access rights to each name space object as the object's owner.

Can change UID to self if has Control Perm

If checked, then when a user has control access to an object, the user can take over ownership of the object by changing the UID of the object to the UID associated with the user.

Can change UID if has Delete Perm on Security ACL

If checked, then when a user is listed with delete permission in the security ACL of the Core Server that owns the object, the user can change the UID of the object to any valid UID.

Object names can contain unprintable characters

If checked, then users of the HPSS system may create files, directories, or other objects with names containing unprintable characters as viewed through the 7-bit ASCII character set. If this option is off, any attempt to name an object using unprintable characters will be disallowed. The range of printable characters is 0x20 (blank) through 0x7E (tilde) in the 7-bit ASCII character set.

Trashcan settings

Trashcans Enabled

If checked, the HPSS trashcan feature will be enabled after the Core Servers are recycled or reinitialized.

Number of Trashcan Threads

Specifies the number of background threads that are to be employed deleting entries from trashcans. This is a performance-tuning knob. For example, if the number of entries in

trashcans seems to be high, or if there is a need to reclaim disk space, or if there is some other trashcan-related need, the administrator may wish to increase the number of trashcan-related background threads so that these tasks can be performed more quickly. On the other hand, if the number of trashcan background threads is taking up resources that might be needed elsewhere, the administrator may wish to decrease the number of threads.

Trash Incinerator Interval

The trashcan deletion threads run periodically in the background. This knob is provided to control the periodicity. **Trash Incinerator Interval** is specified in seconds. For example, if this setting is 3600 the trashcan deletion threads will kick off every hour.

Trash Unlink Eligible Time

The amount of time, in seconds, that an entry will remain in the trashcan before it is eligible for permanent deletion by the garbage collector. For example, if **Trash Unlink Eligible Time** is set to 691200 then entries would be allowed to remain in trashcans for up to 8 days. The default setting is 864000 (10 days).

Trash Statistics Interval

The thread that gathers trashcan statistics runs periodically to gather statistics about all entries that are currently in trashcans. This knob controls how often this thread runs and is specified in seconds. For example, if **Trash Statistics Interval** is set to 7200 the trashcan statistics thread would run every two hours.



Advice

*The administrator should be aware that the trashcan statistics thread could consume significant system resources when run on a system having a large number (millions) of entries in trashcans. In such a case the administrator might wish to increase the **Trash Statistics Interval** value so that the thread runs less often.*



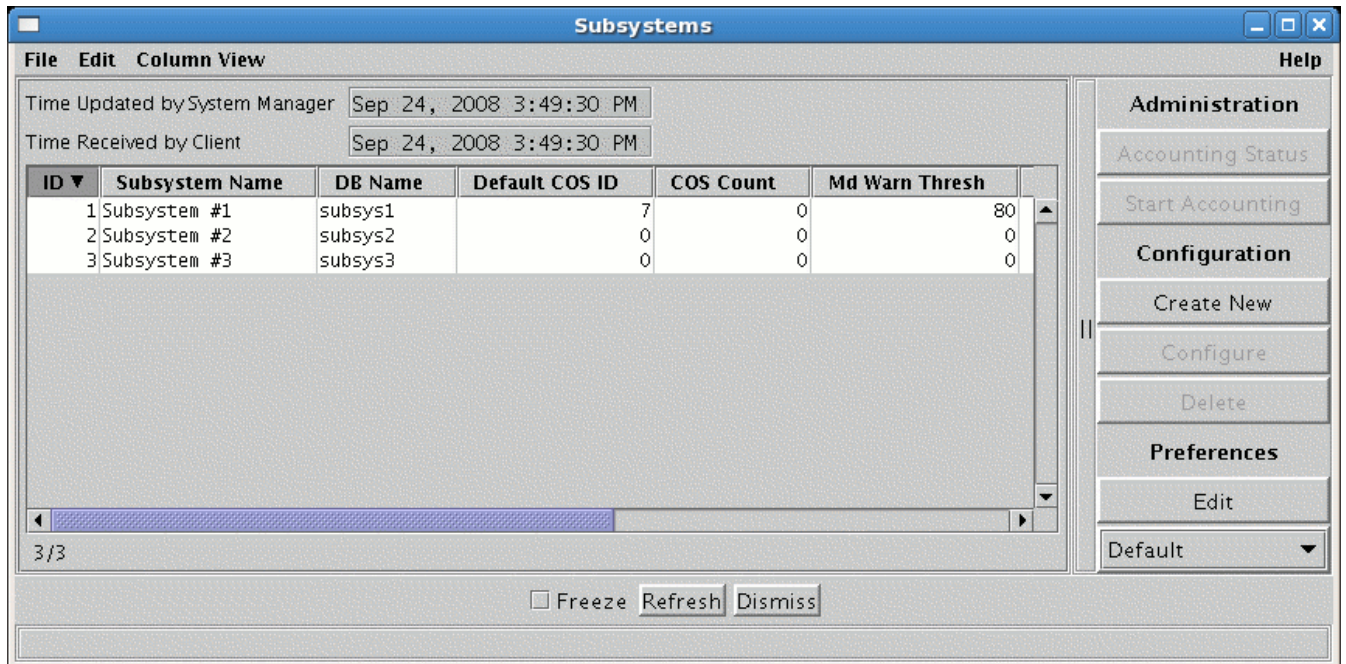
*When changes are made to trashcan settings, the Core Servers must be reinitialized or restarted for the new settings to take effect. It is recommended to use the **Apply Config** button on the HPSS Health and Status window to reinitialize the server(s).*

4.2. Storage subsystems

Every HPSS system consists of at least one subsystem. This chapter provides instructions and supporting information for creating, modifying, and deleting them.

For a conceptual overview of subsystems, refer to the sections *Storage subsystems* and *HPSS storage subsystems* in the *HPSS Installation Guide*, both under the *HPSS basics* chapter.

4.2.1. Subsystems list window



This window lists all the subsystems in the HPSS system and provides the ability to manage these subsystems. To open the window, from the *Health and Status* window select the **Configure** menu, and from there select the **Subsystems** menu item.

To create a new subsystem, click on the **Create New** button. To configure an existing subsystem, select it from the list and click on the **Configure** button. When creating or configuring a subsystem, the *Storage Subsystem Configuration* window will appear.

To delete an existing subsystem, select it from the list and click on the **Delete** button.

Field descriptions.

Subsystem List table

For details on the columns listed for each subsystem, refer to *Section 4.2.3, “Storage Subsystem Configuration window”*.

Administration buttons.

Accounting Status

Opens the *Accounting Status* window which displays the status and statistics from the last accounting run. See *Section 14.2.2.1, “Generating an accounting report”* for more information. You must first select a configured subsystem by clicking on the subsystem entry from the list of displayed subsystems.

Start Accounting

Start the Accounting utility to generate an accounting report. See the *Accounting Status* window for updates from the utility. You must first select a subsystem by clicking on the subsystem entry from the list of displayed subsystems.

Configuration buttons.

Create New

Create a new storage subsystem definition by opening a *Storage Subsystem* window containing default values for a new subsystem.

Configure

Open the selected subsystem configuration for editing. This button is disabled unless a subsystem is selected in the list.

Delete

Delete the selected subsystems. This button is disabled unless a subsystem is selected in the list.



Always contact HPSS support before deleting a subsystem definition. An improperly deleted subsystem can cause serious problems for an HPSS system. Refer to Section 4.2.5, “Deleting a storage subsystem”.

Related information.

Section *Storage subsystems* of the *HPSS Installation Guide* (in chapter *HPSS basics*)

Section *HPSS storage subsystems* of the *HPSS Installation Guide* (in chapter *HPSS basics*)

Section 14.2.2.1, “Generating an accounting report”

4.2.2. Creating a new storage subsystem

The steps detailed in the following sections should be followed to create a new storage subsystem within an HPSS system. These steps are:

1. Create the storage subsystem metadata (*Section 4.2.3.1, “Create storage subsystem metadata”*)
2. Create the storage subsystem configuration (*Section 4.2.3.2, “Create storage subsystem configuration”*)
3. Create the storage subsystem servers (*Section 4.2.3.3, “Create storage subsystem servers”*)
4. Assign a Gatekeeper to the subsystem, if required (*Section 4.2.3.4, “Assign a Gatekeeper if required”*)
5. Assign storage resources to the storage subsystem (*Section 4.2.3.5, “Assign storage resources to the storage subsystem”*)
6. Create the storage subsystem fileset and junction (*Section 4.2.3.6, “Create storage subsystem fileset and junction”*)
7. Configure migration and purge policy overrides (*Section 4.2.3.7, “Migration and purge policy overrides”*)
8. Configure storage class threshold overrides (*Section 4.2.3.8, “Storage class threshold overrides”*)

4.2.3. Storage Subsystem Configuration window

Storage Subsystem Configuration

File Edit Help

Subsystem ID: 1

Default Class of Service: NONE

Subsystem Name: Subsystem #1

Database Name: subsys1

Gatekeeper: NONE

ID	COS Name	Default COS	Allow
1	Class of Service 1	yes	<input checked="" type="checkbox"/>
2	Class of Service 2 (Tape)		<input checked="" type="checkbox"/>
3	Class of Service 3		<input checked="" type="checkbox"/>
4	Class of Service 4		<input checked="" type="checkbox"/>

Metadata Space Warning Threshold: 0 percent

Metadata Space Critical Threshold: 0 percent

Metadata Space Monitor Interval: 0 seconds

DB Log Monitor Interval: 0 seconds

COS Change Stream Count: 0

Update Delete Start Over Start New Clone (full) Dismiss

This window allows an administrator to manage the configuration of a storage subsystem.

The **Add** button is only displayed during the creation of a new configuration. The **Update** button is displayed when an existing configuration is being modified.

To open this window for creation of a new subsystem, click the **Create New** button on the *Subsystems* window. To open this window for an existing subsystem, select the subsystem from the *Subsystems* window and click the **Configure** button.

Field descriptions.

The following fields may be set upon initial creation of the subsystem or modified as configuration needs change.



After making any changes to the following fields, the Core Server for this subsystem will need to be reinitialized or restarted to load the new configuration. It is recommended to

use the **Apply Config** button on the HPSS Health and Status window to reinitialize the server(s).

Subsystem ID

A unique, positive integer ID for the storage subsystem. This field may only be set at create time. The default value is the last configured subsystem ID number plus "1". The default subsystem ID can be overwritten but if the new number already exists, an attempt to save the configuration will fail.

Default Class of Service

This value overrides the global Default Class of Service specified on the *Global Configuration* window. (Refer to Section 4.1, "Global Configuration window".) This default Class of Service applies to this storage subsystem only. The default value is `None`.



Advice

HPSS Core Servers use a default Class of Service to store newly-created files when the HPSS user does not specify a COS or any hints with the creation request. The global configuration specifies a default COS for an entire HPSS installation. Selecting a COS on the Storage Subsystem Configuration window allows the global value to be overridden for a particular subsystem.

If the field is blank, the global default COS will be used. If no Classes of Service are configured, this value can be updated after the Classes of Service are in place.

Subsystem Name

The descriptive name of the storage subsystem. This field may be set only when the storage subsystem is created. The name should be unique and informative. It can contain a character string up to 31 bytes long. The default value is "Subsystem #<ID>".

Database Name

The name of the database to be used to store the metadata for the storage subsystem.

Gatekeeper

The default value is `None`.



Advice

If an appropriate Gatekeeper has not yet been configured, simply leave this configuration entry blank. It can be updated after the Gatekeeper is in place.

Allowed COS list

A list of Classes of Service that can be used by this subsystem. To allow a COS to be used by this subsystem, the corresponding check box must be selected in the Allow column of the list. At least one COS must always be selected. The user will not be permitted to deselect the COS defined to be the Default Class of Service. If this subsystem configuration does not have a Default Class of Service defined, then the COS chosen as the global configuration's Default Class of Service cannot be deselected.



*Note that a newly created COS will not appear in the selection list until the changes have been applied by clicking the **Apply Config** button on the HPSS Health and Status window or the Core Server and Migration/Purge Server associated with the subsystem*

have been reinitialized or recycled. When new Classes of Service are added, the initial allowed state for that COS is determined by the current setting for the other Classes of Service. If all previous Classes of Service were allowed, the new COS will be allowed; otherwise, the new COS will be disallowed.



Advice

By default, the servers in a subsystem are able to use any configured COS. This table allows an administrator to prevent a subsystem from using particular Classes of Service.

When a new Class of Service is added to a system, it will automatically be enabled for all subsystems which have no disabled Classes of Service. It will be disabled in all other subsystems. If this is not the desired configuration, the COS will have to be allowed or disallowed for each subsystem individually.

Disallowing all Classes of Service in a subsystem is not permitted.

Metadata Space Warning Threshold

The Core Server in this subsystem will issue a warning alarm periodically and set its Opstate to Major when the percentage of used space in any DB2 tablespace in this subsystem's database exceeds this value.

Metadata Space Critical Threshold

The Core Server in this subsystem will issue a critical alarm periodically and set its Opstate to Critical when the percentage of used space in any DB2 tablespace in this subsystem's database exceeds this value.

Metadata Space Monitor Interval

The Core Server for this subsystem will check the metadata usage statistics at the indicated interval, specified in seconds. If a value of "0" is specified, the global configuration setting will be used for this storage subsystem. The minimum value for this field is 1800 seconds (30 minutes).

DB Log Monitor Interval

The Core Server will check the consistency of Database Logs and Backup Logs at the indicated interval, specified in seconds. The logs are consistent if both primary and backup log directories exist and contain log files with the same names. The minimum value for this field is 300 seconds (five minutes). A value of "0" will turn off DB Log monitoring. This field may be overridden on the storage subsystem configuration.

COS Change Stream Count

The number of background threads that run in the Core Server to process Class of Service change requests. If a value of "0" is specified, the global configuration setting will be used for this storage subsystem.

4.2.3.1. Create storage subsystem metadata

Before creating the subsystem metadata, you must have created the subsystem database that will be used by the subsystem and have created appropriate tablespaces for the database.

You should review and perform the steps in the following sections in the *HPSS Installation Guide* before allocating metadata space for the new subsystem:

- *HPSS storage subsystems* (in chapter *HPSS basics*)
- *HPSS infrastructure storage space* (in chapter *HPSS planning*)
- *HPSS file systems* (in chapter *HPSS planning*)
- *HPSS metadata space* (in chapter *HPSS planning*)
- *Configure HPSS - Secondary subsystem machine* (in chapter *HPSS installation and infrastructure configuration*)

4.2.3.2. Create storage subsystem configuration

Use the following steps to create a storage subsystem configuration:

1. Decide on a unique descriptive name for the new storage subsystem. SSM will automatically choose the name "Subsystem #N", where N is the subsystem ID selected by SSM. The name of the new storage subsystem may be changed by the administrator at subsystem configuration time only.
2. From the *HPSS Health and Status* window, select **Configure** from the menu. Then select **Subsystems**. You will then be presented a window that lists currently configured subsystems. Select the **Create New** button to bring up another window to configure the new subsystem.
3. Enter the name you have chosen for the subsystem in the **Subsystem Name** field and enter the ID of the subsystem in the **Subsystem ID** field. Enter the name of the database you have chosen to contain the tables for this subsystem in the **Database Name** field.
4. Decide which Classes of Service are to be supported by the new storage subsystem. SSM will automatically select all Classes of Service to be supported, but the administrator can modify these choices at any time. Also, decide on a default Class of Service for this storage subsystem. By default SSM leaves this field blank, which means that the default Class of Service specified in the global configuration will apply to the new storage subsystem as well. The administrator may choose to override the global value by using the subsystem configuration value at any time.
5. Decide whether gatekeeping or account validation are needed for this storage subsystem. If either is required, a Gatekeeper will need to be configured for this storage subsystem. If the required Gatekeeper is already configured, simply add it to your storage subsystem's configuration. However, if it is not yet configured, it will be necessary to wait until *Section 4.2.3.4, "Assign a Gatekeeper if required"* to add the Gatekeeper.
6. Set the metadata space thresholds and the update interval. Typical values are "75" for warning, "90" for critical and "300" to have the metadata space usage checked every 300 seconds.
7. Set the **DB Log Monitor Interval**. The minimum value is 300 seconds; the typical value is 1800 seconds.
8. Click the **Add** button to store the configuration.

4.2.3.3. Create storage subsystem servers

The new storage subsystem must contain a single Core Server. If migration and purge services will be needed then a Migration/Purge Server is also required. See the following sections for instructions on configuring these new servers:

- *Section 5.2, “Server configuration”*
- *Section 5.2.2, “Core Server-specific configuration”*
- *Section 5.2.5, “Migration/Purge Server (MPS)-specific configuration”*

On each server’s basic configuration window, be sure to assign the server to the new storage subsystem.

4.2.3.4. Assign a Gatekeeper if required

If gatekeeping or account validation are needed and an appropriate Gatekeeper is not already configured, a new Gatekeeper will be required. See the *Gatekeeper* section of the *HPSS Installation Guide* (in chapter *HPSS Planning*), and *Section 5.2.3, “Gatekeeper-specific configuration”* for instructions on configuring the Gatekeeper.

Be sure to assign this Gatekeeper to the appropriate storage subsystem by choosing it from the Gatekeeper selection list on the appropriate Storage Subsystem Configuration window.

4.2.3.5. Assign storage resources to the storage subsystem

Storage resources are tied to a storage subsystem through storage classes. To determine which storage classes belong to a particular subsystem, look up the configuration for each Class of Service available to the storage subsystem. Each Class of Service is tied to one storage hierarchy which is in turn tied to some number of storage classes. Storage resources must then be assigned to the storage classes which are used in the referenced hierarchies. See *Section 8.1.2, “Creating storage resources”* for details.

4.2.3.6. Create storage subsystem fileset and junction

The HPSS name space consists of filesets that are connected by junctions. The top (root directory) of the HPSS name space is the RootOfRoots fileset managed by the root Core Server. The rest of the name space is built by pasting filesets together using junctions.

Since each subsystem has its own Core Server, it also has its own root fileset. This fileset needs to be connected to the HPSS name space in order for files to be stored in the new subsystem. This is accomplished by creating a junction from the HPSS name space to the root fileset of the new subsystem.

See the *Chapter 10, Filesets and junctions* chapter for more information.

4.2.3.7. Migration and purge policy overrides

The migration and purge policies contain two elements, the basic policy and the storage subsystem-specific policies. This can be seen on the *Migration Policy* and *Purge Policy* windows. If a given migration or purge policy does not contain any subsystem-specific policies, then the basic policy applies across all storage subsystems and no other configuration is needed. If it is desired for migration or purge to behave differently than specified in the basic policy in a given storage subsystem, then a storage subsystem-specific policy should be created for that subsystem. A subsystem-specific policy allows some or all of the values in the basic policy to be overridden for the given storage subsystem. See *Section 6.4.2, “Migration policy configuration”* and *Section 6.5.2, “Purge Policy configuration”* for more information.

4.2.3.8. Storage class threshold overrides

The warning and critical thresholds given on the *Storage Class Configuration* window apply across all storage subsystems unless specified otherwise. The **Subsystem Thresholds** button on the *Configured Storage Class* list window allows the default thresholds to be overridden for specified storage subsystems. See *Section 6.1.1, “Configured Storage Classes window”* for more information.

4.2.4. Modifying a storage subsystem

If modifications are made to an existing storage subsystem configuration, the Core Server and the Migration/Purge Server for the subsystem must be recycled.

4.2.5. Deleting a storage subsystem



*Always contact HPSS support before deleting a subsystem definition. An improperly deleted subsystem can cause serious problems for an HPSS system.
It is critical that no files or directories exist within a storage subsystem before it is deleted.
It is important to verify that all of the DB2 metadata tables associated with the storage subsystem being deleted are empty.*

To verify that all files have been removed from the subsystem, perform the following steps:

1. Run the **dump_acct_sum** utility on the subsystem. Be sure to specify the subsystem ID with the **-s** option as this utility defaults to subsystem 1. The output from the utility should indicate that zero files exist in the subsystem.
2. As a second verification do the following:
 - a. Run the **db2** command line utility program.
 - b. Connect to the subsystem database (for example, connect to `subsys1`).
 - c. Set the schema (for example, set schema `hpss`).
 - d. Issue the following SQL command:

```
db2> select count(*) from bitfile
```

The result of the command should indicate "0" rows in this table.

- e. Issue the following SQL command:

```
db2> select count(*) from nsubject
```

The result of the command should indicate "2" rows in this table.

3. If any of these checks gives an unexpected result, do not delete the subsystem. Contact HPSS support.



*When deleting an existing storage subsystem, it is critical that all of the different configuration metadata entries described in the *Section 4.2.3, “Storage Subsystem**

Configuration window” for the storage subsystem be deleted. If this is not done, configuration metadata entries associated with the subsystem will become orphaned when the subsystem configuration itself is deleted. This situation is difficult to correct after the subsystem is deleted.

Once it has been verified that the storage subsystem contains no files or directories, deleting the subsystem may be accomplished by reversing the steps used to create the subsystem. These steps are listed in the order which they should be performed below. For information on each step, see the corresponding instructions under *Section 4.2.2, “Creating a new storage subsystem”*.

1. Delete all storage resources assigned to the subsystem’s Core Server. See *Section 8.2.1, “Deleting storage resources”*.
2. Delete all filesets and junctions used to link the storage subsystem into the HPSS name space. See *Section 10.6, “Deleting a junction”*.
3. Delete all storage class subsystem-specific thresholds associated with the storage subsystem being deleted. This is done by setting these thresholds back to the storage class "default" thresholds.
4. Delete all subsystem-specific migration and purge policy entries associated with the storage subsystem being deleted. See *Section 6.4.2.4, “Deleting a migration policy”* and *Section 6.5.4, “Deleting a purge policy”* for more information.
5. Delete all Core Servers, the Migration/Purge Server, and the DMAP Gateway Server residing within the storage subsystem being deleted. See *Section 5.2.9, “Deleting a server configuration”*.
6. Delete the subsystem configuration.
7. Remove the database associated with the subsystem.

4.3. HPSS trashcans

The HPSS trashcan feature is similar to the trashcan feature found on many other systems. The purpose of all trashcan implementations is to provide a method to recover from deleting an entry that is still needed.

On and off

A notable difference from other trashcan implementations is HPSS’s ability to turn the trashcan feature on and off globally. Sites not wishing to run with trashcans can easily configure the feature to be off. This is done by unchecking the Trashcans Enabled box in the Trashcan Settings section of the *Global Configuration* window. (Refer to *Section 4.1, “Global Configuration window”*.) The default HPSS configuration is to have trashcans turned off.

Sites choosing to run with trashcans turned off will not have the ability to recover deleted entries. With trashcans turned off, all deleted entries will be instantaneously and irrevocably lost.

When the trashcan feature is turned on, all deleted entries are moved into the `.Trash` directory for each user, where the entries will remain for a configurable amount of time. While entries are in the `.Trash` directory they can be recovered.

4.3.1. The trashcan feature is on

The remainder of this trashcan discussion and description assumes that the trashcan feature has been turned on.

Creating a .Trash directory

When an entry is deleted HPSS implicitly creates the trash directory (if the trash directory doesn't already exist), usually in the user's home directory. There is no need to explicitly create the .Trash directory.

If an entry is deleted from the root fileset, the .Trash directory is created in the user's home directory. The trash directory will be named:

```
~/ .Trash
```

However if the entry being deleted does not reside in the root fileset, the trash directory will be created (if it doesn't already exist) in the root node of the fileset that contains the entry, and the trash directory will be named:

```
./.Trash/<user name>
```

The reason for putting entries into .Trash directories in the root nodes of filesets is because HPSS does not allow objects to be moved across fileset boundaries. This restriction is what prevents us from moving the entry into the .Trash directory in the user's home directory.

If user Joe Smith deletes `file6` from a fileset (other than the root fileset) that is accessed via the junction `JunctionToFS`,

```
rm /home/smith765/JunctionToFS/project3/file6
```

`file6` will be moved into the .Trash of the root node of the fileset as follows:

```
./.Trash/smith765/file6.<ObjectId>
```

Names

When entries are deleted they are moved into the .Trash directory. The .Trash directory is a flat, one-level directory. To prevent name collisions, the name of every object is made unique by appending the object's `ObjectId` to the end of its name. (Every HPSS object is assigned a unique number known as its `ObjectId`).

If user Joe Smith deletes file `Foo77` (which has `ObjectId` 54321) from his root fileset, it is moved into the .Trash directory with the following name

```
~/ .Trash/Foo77.54321
```

What can be deleted?

Files, directories, junctions, and links can be put in the trash. Note that, as usual, the root nodes of filesets cannot be deleted by clients. Further note that, as always, directories must be empty before they can be deleted.

Immediate deletion

To cause an object to be immediately deleted, a user can delete a previously deleted object from the `.Trash` directory. There is no single-step permanent deletion method when the trashcan feature is enabled.

If user Joe Smith wants file `Foo77` to be permanently deleted immediately, this can be accomplished by deleting `Foo77`, then deleting it again from the trashcan with its appended `ObjectId`. If assistance is needed in obtaining the `ObjectId`, contact HPSS support.

```
rm ~/OriginalPathTo/Foo77
rm ~/.Trash/Foo77.54321
```

Object recovery

If user Joe Smith needs to recover `Foo77`, Joe can **cd** to the `.Trash` directory and move it out of the `.Trash`.

```
mv ./Foo77.54321 ../SafeDirectory
```

Another way to recover

In addition to moving an entry to recover it, there may be some interfaces that provide an undelete command which will move objects from the `.Trash` directory to their original location, assuming the original location still exists.

Finding entries in the trash

Because `.Trash` directories are single level directories, finding the desired entry among many similarly named `.Trash` entries may prove to be difficult. For example, the `.Trash` directory could contain the following versions of the file `Foo77`:

```
Foo77.54321
Foo77.834239
Foo77.345678
```

To determine which of these files was the last one deleted, list the contents of the `.Trash` directory and review the time created. (Not all applications may support the flags shown in this example.)

```
ls -lt
```

Trash Time

The length of time that an object will remain in the trashcan before being permanently deleted is configurable and will probably be different at each HPSS site. The default length of time set by HPSS is 10 days. Note that this length of time is applied globally to all trashcans in the HPSS system. This length of time can be changed by an HPSS system administrator via the SSM's *Global Configuration* screen.

Note that the instant an object is deleted, its Trash Time starts counting. And whether the Trash Time is 60 seconds or 60 days, when the Trash Time is up, the object becomes eligible for permanent deletion and will likely be permanently deleted relatively quickly after becoming eligible, assuming the permanent deletion threads are able to perform as expected.

4.3.2. Trashcan operations for the administrator

There are some trashcan operations that can only be performed by the HPSS administrator. All of these trashcan configuration operations can be performed on the SSM's *Global Configuration* screen. Note that various trashcan-related counts and statistics can be viewed from the *Core Server Information* window.

TrashEnabled - enabling the trashcan feature

In the *Global Configuration* window, there is a check box for enabling the trashcan. Checking the box enables the trashcan, unchecking the box disables it. The box is unchecked by default. (Refer to *Section 4.1, "Global Configuration window"*.)

Trash Unlink Eligible Time

Entries in trashcans are permanently deleted only after a specified amount of time has elapsed. **Trash Unlink Eligible Time** is used to specify that amount of time in seconds. For example, if **Trash Unlink Eligible Time** were to be set to "691200" then entries would be allowed to remain in trashcans for at least 8 days. The default setting is 864000 seconds which is 10 days.

Number of Trashcan Threads

Specifies the number of background threads that are to be employed permanently deleting entries from trashcans. This is a performance-tuning knob — for example, if the number of entries in trashcans seems to be high, or if there is a need to reclaim space consumed by trash entries, or if there is some other trashcan-related need, the administrator may wish to increase the number of trashcan-related background threads so that these tasks can be performed more quickly. On the other hand, if the number of trashcan background threads is taking up resources that might be needed elsewhere, the administrator may wish to decrease the number of threads.

If **Number of Trashcan Threads** is set to zero, HPSS interprets this as a request to no longer delete entries from `.Trash` (the NSTRASH database table). If the **Number of Trashcan Threads** is zero, deletions will continue to move entries into `.Trash` directories, but these entries will not be deleted until the **Number of Trashcan Threads** is set to a positive nonzero integer value.

Trash Incinerator Interval

The trashcan deletion threads run periodically in the background. This knob is provided to control this periodicity. **Trash Incinerator Interval** is specified in seconds. If **Trash Incinerator Interval** is set to "3600", the trashcan deletion threads will kick off every hour.

Trash Statistics Interval

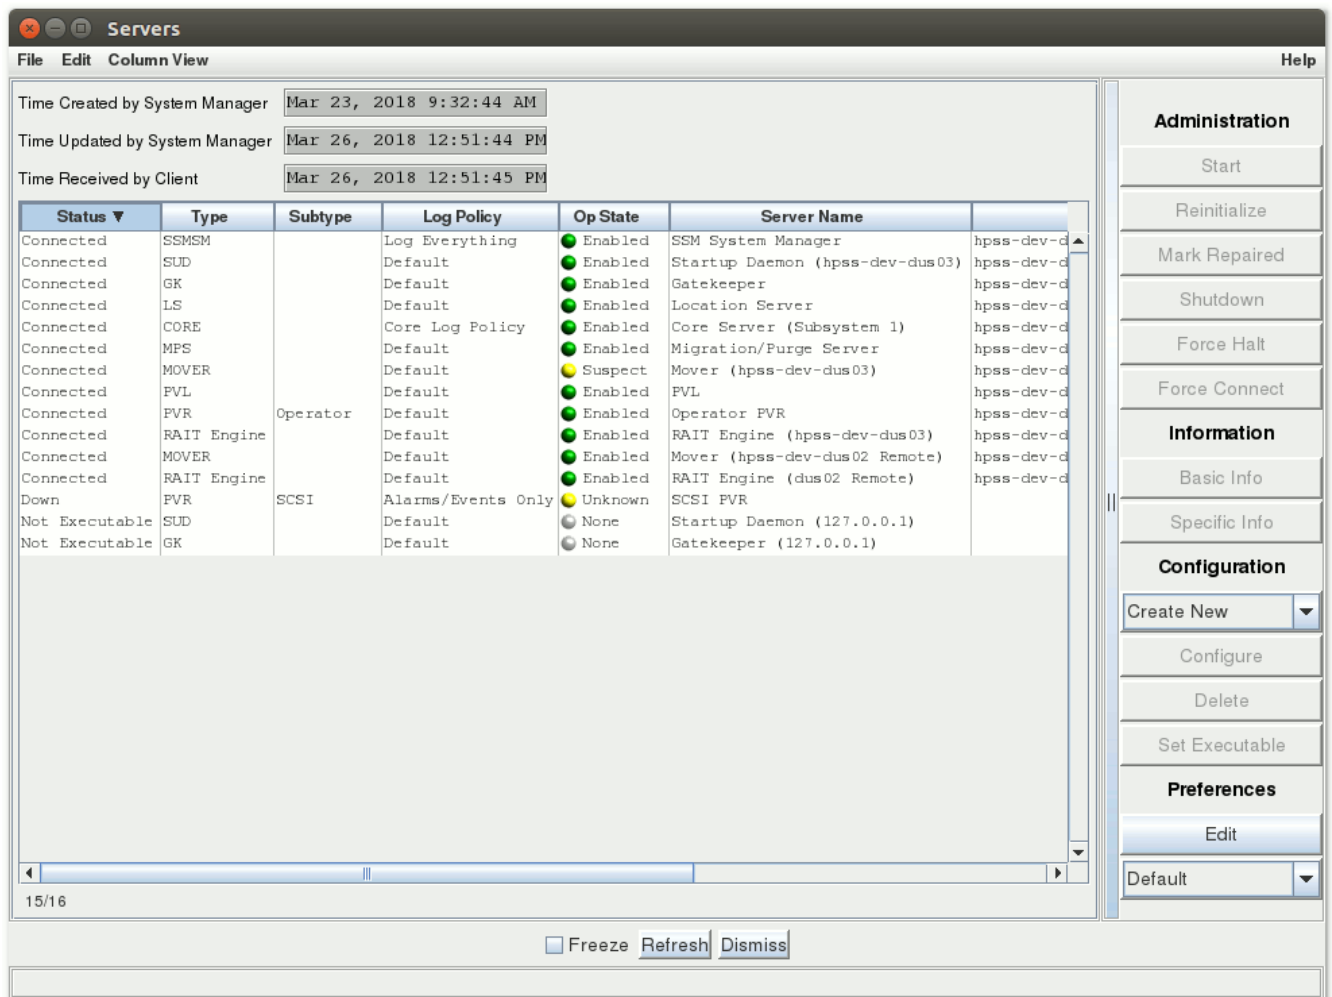
The thread that gathers trashcan statistics runs periodically to gather statistics about all entries that are currently in trashcans. This knob controls how often this thread runs, and is specified in seconds. For example, if **Trash Statistics Interval** were to be set to "7200", the trashcan statistics thread would run every two hours. The administrator should be aware that the trashcan statistics thread might consume significant system resources when run on a system having a large number (millions) of entries in trashcans. In such cases the administrator might wish to increase the **Trash Statistics Interval** value

so that the thread runs less often. Also, note that this thread will continue to run if the trashcan is not enabled.

Chapter 5. HPSS servers

Most HPSS server administration is performed from the SSM graphical user interface *Servers* list window. Each HPSS server has an entry in this list.

5.1. Server list



This window facilitates management of the configured HPSS servers. From this window, an HPSS server can be started, shut down, halted, reinitialized, and notified of repair. Once a server is up and running, SSM monitors and reports the server state and status. Information on running servers may also be viewed and updated via this window.

If multiple entries are selected in the server list when a server operation is invoked, the operation will be applied to all selected servers (although some servers don't support all operations).

The server entries displayed on the window can be sorted by each column category. To sort the server entries by status, for example, click on the Status column title. The actual display can vary greatly by the setting of window preferences. The Column View menu item can be used to select which

columns of the table are visible, and Preferences can be used to further refine and filter which servers are visible. Preferences can also be saved and automatically reloaded by SSM when new sessions are started.

At times, the server list may update quite frequently with new Status or Op State information. If you select either of these columns for sorting, servers are likely to change their positions in the list rapidly. Sometimes this makes the list hard to use, in which case you should consider selecting a more static column for sorting or check the **Freeze** button to keep the list from updating.

Field descriptions.

Server List

The main portion of the window which displays various information about each server.

ID

A unique numerical value assigned by SSM when each server starts. This value can change each time SSM is restarted. In the default preferences, this column is not shown.

Status

The server execution and connection status as determined by SSM. The reported status will be one of the following:

- `Connected` when the Server is up and running and communicating normally with SSM.
- `Up/Unconnected` when the Server is up and running (according to the Startup Daemon) but SSM cannot connect to it. Server cannot be completely controlled and monitored through SSM.
- `Down` when the Server is down. SSM can be used to start the server.
- `Indeterminate` when the server's state cannot be determined by SSM and the Startup Daemon is either not running or not connected to SSM.
- `Check Config` when SSM detected an incomplete or inconsistent configuration for the server. The server's configuration should be carefully reviewed to ensure that it is correct and complete. Check the *Alarms and Events* window and the HPSS log file to view SSM alarm messages related to configuration problems. This situation can be caused by:
 - A DB2 record required by the server is missing or inaccessible.
 - The principal name configured for the server does not match the `HPSS_PRINCIPAL_*` environment variable for the server's type.
- `Not Executable` when the server is configured as non-executable. SSM will not monitor the server's status.

Note: the server could still be running. It may have been started outside the control of SSM, or it may have been running when its executable flag was changed.

- `Deleted` when the server configuration has been deleted. In the default preferences, this Status type is filtered off.

Note: deleted servers will be removed from the list when SSM is restarted.

In addition to the above status values, the **Status** field also reports the transient status for the server as the result of the user request on the server as follows:

- Starting... when the server is being started.
- Stopping... when the server is being shut down gracefully.
- Halting... when the server is being forcibly halted.
- Reiniting... when the server is reinitializing.
- Connecting... when SSM is trying to establish a connection to the server.
- Repairing... when the server is repairing its states and statuses.

A server that is configured to execute and is running should have a `Connected` status. If its status is anything other than `Connected` (excluding the transient status values), one of the following actions should be taken:

- If the server status is `Up/Unconnected`: Monitor the server status closely for a few minutes. SSM will periodically try to establish a connection with the server. The **Force Connect** button can be used to speed this process.
- If the server status is `Down`: Use the **Start** button to start the server. The Startup Daemon will ensure that only one instance of the server is running.
- If the server status is `Indeterminate`: Verify whether the server is running using an operating system command such as `ps`. If the server is not running, start it. If the server is running, ensure that the Startup Daemon configured for the same node is running and has a connection to SSM. If the Startup Daemon is not running, start it using the `/opt/hpss/bin/rc.hpss` script on the appropriate node. Otherwise, use the **Force Connect** button to establish the connections for the server and the Startup Daemon. If this does not correct the server's status, review the *Alarms and Events* window to search for problems that the server and SSM may have reported. In addition, review the HPSS logs for the server's and SSM's log messages to help determine the problems.

SSM periodically polls the execution and connection status of the servers and updates the **Status** fields when there are changes. The rate of these updates depends on the client's refresh rate (see the `hpssgui` or `hpssadm` man pages for more details).



If a server is configured as executable but is not running, SSM will treat it as an error. Therefore, if a server is not intended to run for an extended period, its Executable flag should be unchecked. SSM will stop monitoring the server and will not report the server-not-running condition as a critical error. This will also help reduce the work load for SSM.

Type

Indicates the type of the server in acronym form. Possible values include CORE, GK, LS, MPS, MOVER, PVL, various PVR values, SSMSM, and SUD. See the glossary for the meanings of these acronyms.

Subtype

Indicates the subtype of the server, sometimes in acronym form. Not all servers have subtypes. Possible values (particularly PVR subtypes) include SCSI, STK, Operator, and AML. See the glossary for the meanings of these acronyms.

Subsystem

The name of the storage subsystem to which the server is assigned. Only servers of type Core and MPS should have a subsystem name. For all other server types this column should be blank.

Op State

The operational state of the server as reported by the server itself:

- `Enabled` indicates the server is operating normally.
- `Disabled` indicates the server is not operational, usually due to a shutdown or halt request.
- `Suspect` indicates the server may have a problem.
- `Minor` indicates the server encountered a problem that does not seriously affect the HPSS operation.
- `Major` indicates the server encountered a problem that may seriously impact the overall HPSS operation.
- `Broken` indicates the server encountered a critical error and shut itself down.
- `Unknown` indicates the server's state is unknown to SSM. SSM should be able to obtain an Operational State for the server but cannot because SSM cannot communicate with it, and the server has not set its state to `Disabled` or `Broken`. The server may or may not be running. The reason may simply be that the server is marked executable but is not running.
- `Invalid` indicates SSM has obtained an Operational State from the server, but its value is not recognized as a valid one. This may indicate a bug in the server or a communications problem which is garbling data.
- `None` indicates the server does not have an Operational State (because the server is not executable, for example), or its Operational State is, by design, unobtainable.

SSM will periodically poll the server's operational state and update the **Op State** field in the *Servers* window when there are any changes. The rate of the polling is controlled by the SSM client list refresh rate (see the **hpssgui** or **hpssadm** man pages).

Server Name

The descriptive name of the server.

Host

The name of the host on which the server is running. The **Host** is obtained by looking at the RPCs that are sent from the server back to the System Manager. If the server is not running, then this field will be blank. In most cases, this field will match the **Execute Host**; sometimes, in the case of multi-homed systems, this field won't match the **Execute Host**.

Execute Host

The **Execute Hostname** field from the server's basic configuration record. This field is intended to specify the hostname on which the server is supposed to run; however, no checking is done to verify if the server is actually running on the specified host. This field is only used by the SSM to locate the Startup Daemon that manages this server. The field displayed must match exactly the **Execute Hostname** field specified in the configuration record of the startup daemon that is to manage this server.

UUID

The universal unique identifier for the server. In the default preferences, this column is not shown.

Administration buttons.

This group of buttons allows you to perform administrative operations on the selected servers. All of the buttons are disabled unless one or more servers are selected and the operation is applicable to at least one of the selected servers.

After clicking one of these buttons, you will usually be prompted to confirm your request before it is carried out. If you make an obviously invalid request, such as asking to start a server which is not executable, SSM will tell you about it, but otherwise no harm is done. You could, for example, select a range of servers and ask to start them all, knowing that some of the selected servers are already running. Those that are not running will be started, but nothing will be done to those that are already running.

The status bar at the bottom of the window displays a message when the requested operation has begun. For information on the success or failure of the operation, monitor the Status and Op State columns for the server.

Start

Start the selected servers. The System Manager will notify the Startup Daemon to start the selected servers.

Reinitialize

Send a "reinitialize" command to the selected servers. The Startup Daemon does not support reinitialization. If a server does not support reinitialization, the button will not be enabled.

Mark Repaired

Clear displayed error status in the selected servers. Sometimes server states such as the Operational State will continue to indicate an error condition after the cause of the error has been fixed. If so, you can use the **Mark Repaired** button to clear its error states. Note that this does not do anything in hardware or software to actually repair an error; it just asks the server to clear its error condition. If you mark a server repaired when it still has a problem, the error states will be cleared but may quickly return.

Shutdown

Command the selected servers to shut down. This command should be the first you use to shut down servers. After issuing this command, wait a minute or two for the server to complete the shutdown process. Some servers shut down very quickly after receiving the command, others, particularly the Core Server, may require two minutes or more to shut down. During this time the server is attempting to finish work it has started, while rejecting new work. Be patient; watch the *Alarm and Events* window for messages indicating the server has terminated.

Force Halt

Command the selected servers to stop immediately. This should be done only if a shutdown request has failed to stop the servers, or if the intention is to shut servers down as quickly as possible. This request will cause a SIGKILL signal to be sent to the selected servers. This command is meant to be used as a last resort if a server hangs up or otherwise won't respond to the Shutdown command.

Force Connect

Request the System Manager to immediately attempt to connect to the selected servers. The System Manager routinely attempts to connect to any unconnected servers; using this button will simply cause the next attempt to occur right away, instead of after the normal retry delay.

Information buttons.

These buttons allow you to open information windows for servers. They will be enabled only when the selected server supports a particular information window.

Since the requested information is obtained by calling the selected server, the server must normally have a Connected status for the request to succeed.

Basic Info

Opens the *Basic Server Information* window for the selected server.

Specific Info

Opens the type-specific server information window for the selected server.

Configuration buttons.

These buttons allow you to start server configuration tasks.

Create New

Allows an administrator to configure a new server by selecting a server type and then filling in a new server configuration. This control is actually not a button, but a drop-down list, allowing the administrator to select the type of new server to be created. It is always enabled.

Configure

Opens a configuration window for each selected server.

Delete

Deletes the selected servers from the system. A confirmation dialog will appear to confirm the action. Before deleting a server, see the warnings and considerations in *Section 5.2.9, "Deleting a server configuration"*.

Set (Non-)Executable

Sets or unsets the Executable configuration flag for each selected server.

5.2. Server configuration

The following HPSS servers should be configured. A Gatekeeper is required only if the site wishes to do gatekeeping or account validation. A RAIT Engine is required only if the site uses RAIT volumes.

See the *Servers* section of the *HPSS Installation Guide* (in the *HPSS basics* chapter) for a description of the purpose of each server:

- Core Server
- Migration/Purge Server
- Gatekeeper (optional)
- Location Server
- Physical Volume Library
- Physical Volume Repository
- Mover
- RAIT Engine (optional)
- Storage System Manager
- Startup Daemon (on each host where an HPSS server will be executing)

The fields of the *Server Configuration* window are divided into the following sections. The Basic Controls section is at the top of the window and the other sections are on individual tabs:

Basic Controls

Server identification and type information.

Execution Controls

Information required to properly control the server's execution.

Interface Controls

Information required for communication between the server and other HPSS servers and clients.

Security Controls

Security settings.

Audit Policy

Object event audit settings.

Specific

Server type-specific settings (only for some server types).

The server-specific configuration section contains configuration parameters specific to a particular server type. Not every type of server has a specific configuration section. The following types of servers have a server-specific configuration section in the *Server Configuration* window:

- Core Server
- Gatekeeper
- Migration/Purge Server

- Mover
- Physical Volume Repository
- RAIT Engine

Although the Location Server does not have a server-specific configuration section in the *Server Configuration* window, there are additional configuration steps outside of SSM that are necessary whenever a Location Server is added or modified.

Details about the specific configuration section and the additional configuration required for each of these server types are described in:

- *Section 5.2.2, “Core Server-specific configuration”*
- *Section 5.2.3, “Gatekeeper-specific configuration”*
- *Section 5.2.4, “Location Server additional configuration”*
- *Section 5.2.5, “Migration/Purge Server (MPS)-specific configuration”*
- *Section 5.2.6, “Mover-specific configuration”*
- *Section 5.2.7, “Physical Volume Repository (PVR)-specific configuration”*
- *Section 5.2.8, “RAIT Engine-specific configuration”*
- Details about all of the other sections on this window, which apply to all server types, are described in *Section 5.2.1, “Common server configuration”*.

To view the *Server Configuration* window for an existing server, bring up the *Servers* list window and select the desired server. Then click the **Configure** button to bring up the configuration window for that server.

To modify a server configuration, bring up the server’s configuration window, modify the desired fields, and click the **Update** button. Some configuration fields may be set by the administrator only at server creation time. These fields will be editable and settable when creating new servers, but will become display-only fields (not editable or settable) when looking at existing configurations. For the modified configuration data to be recognized by the server, if the setting changed allows, the change should be applied by clicking the **Apply Config** button on the *HPSS Health and Status* screen. Otherwise, the server must be reinitialized (if supported) or restarted.

The Startup Daemon does not support "Apply Config" or reinitialization and must be restarted for configuration changes to take effect.

There are two ways to create a new server:

1. To create a new server using standard defaults as a starting point, open the *Servers* list window, click **Create New** and select a server type. This opens a *Server Configuration* window with default values for the selected server type. Make the desired customizations to the window and click the **Add** button.

2. To create a new server that has attributes similar to those of an existing server, select the existing server, click **Configure** and then select **Clone (partial)**. The **Update** button will be changed to **Add**. Make the desired customizations to the values in the window and click the **Add** button to save the new configuration.

There are two ways to delete a server. Before deleting a server, see the warnings and considerations in *Section 5.2.9, “Deleting a server configuration”*. Then, to delete the server:

1. From the *Servers* list window, select the server and click the **Delete** button, or
2. Bring up the *Server Configuration* window and click the **Delete** button.

In both cases you will be prompted to confirm the deletion of the server.

5.2.1. Common server configuration

These sections of the *Server Configuration* window are common to all servers.

5.2.1.1. Basic controls

The screenshot shows the 'Core Server Configuration' window. The 'Basic Controls' tab is active, displaying fields for Server Name, Server ID, UUID, Server Type, Storage Subsystem, and Log Policy. Below this, the 'Execution Controls' tab is selected, showing fields for Execute Pathname, Execute Hostname, Program Number, Version Number, UNIX Username, and Auto Restart Count, along with checkboxes for Executable and Auto Startup. At the bottom, there are buttons for Update, Delete, Start Over, Clone (partial), and Dismiss.

Basic Controls	
Server Name	Core Server (Subsystem 1)
Server ID	7
UUID	4ca0aad4-8e19-11e5-83b3-005056b7004f
Server Type	CORE
Storage Subsystem	Subsystem #1
Log Policy	Core Log Policy

Execution Controls	
Execute Pathname	/opt/hpss/bin/hpss_core
Execute Hostname	hpss-dev-dus03.ccs.ornl.gov
Program Number	536870917
Version Number	1
UNIX Username	root
Auto Restart Count	3
<input checked="" type="checkbox"/> Executable	<input checked="" type="checkbox"/> Auto Startup

Buttons: Update, Delete, Start Over, Clone (partial), Dismiss

The Basic Controls section of the *Server Configuration* window is common to all servers. In the example window above, the server displayed is a Core Server.

Field descriptions.

Server Name

A unique descriptive name given to the server. Ensure that the **Server Name** is unique. A server's descriptive name should be meaningful to local site administrators and operators, in contrast to the server's corresponding UUID, which has meaning for HPSS. For HPSS systems with multiple subsystems it is very helpful to append the subsystem ID to the **Server Name** of subsystem-specific servers. For instance, "Core Server 1" for the Core Server in subsystem 1.

Server ID

A unique value automatically generated and displayed for servers.

UUID

A universal unique identifier that identifies a server to HPSS. A unique value is automatically generated and displayed for new servers.

Server Type

The type of the HPSS Server.

Server Subtype

The subtype of the selected server. This field is only used by the PVRs to specify the type of PVR (for example, STK).

Storage Subsystem

The name of the HPSS storage subsystem to which this server is assigned. This field is required for the Core Server and Migration/Purge Server. For all other servers this field is not displayed. The following rules apply to this field:

- The storage subsystem configuration must exist before configuring Core Servers or Migration/Purge Servers.
- Core Servers and Migration/Purge Servers must have a valid storage subsystem set in this field before SSM will allow the server configuration to be created.
- No more than one Core Server and one Migration/Purge Server can be assigned to any one subsystem.
- This field can only be modified in Add mode.

Log Policy

The log policy used to specify which log message types will be sent to syslog for this server. See *Section 9.2, "Log policies"* for a description of the *Logging Policies* window, for detailed definitions of each log message type, and for information on the system default logging policy.

5.2.1.2. Execution controls

The screenshot shows the 'Core Server Configuration' window with the 'Execution Controls' tab selected. The window has a menu bar with 'File', 'Edit', and 'Help'. The 'Basic Controls' section includes fields for 'Server Name' (Core Server (Subsystem 1)), 'Server ID' (7), 'UUID' (4ca0aad4-8e19-11e5-83b3-005056b7004f), 'Server Type' (CORE), 'Storage Subsystem' (Subsystem #1), and 'Log Policy' (Core Log Policy). The 'Execution Controls' section includes fields for 'Execute Pathname' (/opt/hpss/bin/hpss_core), 'Execute Hostname' (hpss-dev-dus03.ccs.ornl.gov), 'Program Number' (536870917), 'Version Number' (1), 'UNIX Username' (root), and 'Auto Restart Count' (3). There are also checkboxes for 'Executable' and 'Auto Startup', both of which are checked. At the bottom, there are buttons for 'Update', 'Delete', 'Start Over', 'Clone (partial)', and 'Dismiss'.

The **Execution Controls** section of the *Server Configuration* window is common to all servers. In the example window above, the server displayed is a Core Server.

Field descriptions.

Execute Pathname

The UNIX file system path name to a server's executable image file. This file must reside on the node specified by **Execute Hostname**. Use the full UNIX path name; otherwise, the Startup Daemon will try to start the file out of the current working directory of the Startup Daemon.

Execute Hostname

This is the hostname of the node on which the server will execute. It must match the **Execute Hostname** of the Startup Daemon that is to manage this server. For most servers, setting this field is straightforward, but for remote Movers, this indicates the node on which the Mover administrative interface process runs (not the node where the remote Mover process runs). Note that if the **Execute Hostname** is changed, it is likely that the RPC program number will change as well. If the server affected is the System Manager, the SSM configuration file, `ssm.conf`, must be regenerated or edited.

Program Number

The RPC program number. This value must be unique within the node in which the server runs. The administrator cannot override the default program number value when creating or modifying a server configuration.

Version Number

The RPC version number. The administrator can not override the default version number value when creating or modifying a server configuration.

UNIX Username

The UNIX user name under which the server will run. The name must be registered in the local UNIX authentication database (for example, `/etc/passwd`) or the HPSS local password file. The Startup Daemon will use this user ID when starting the server. If there is no such user, the Startup Daemon will not be able to start the server.

Auto Restart Count

Maximum number of automatic restarts allowed per hour. See *Section 5.6.2.4, “Automatic server restart”* for more details.

Executable

A flag that indicates whether an HPSS server can be executed by SSM.

Auto Startup

A flag that indicates whether SSM should attempt to start the server upon startup of SSM. The Startup Daemon must be started before SSM for this option to work.

5.2.1.3. Interface controls

Core Server Configuration

File Edit Help

Basic Controls

Server Name: Core Server (Subsystem 1)

Server ID: 7

UUID: 4ca0aad4-8e19-11e5-83b3-005056b7004f

Server Type: CORE

Storage Subsystem: Subsystem #1

Log Policy: Core Log Policy

Execution Controls **Interface Controls** Security Controls Audit Policy Specific

Maximum Connections: 200

Thread Pool Size: 200

Request Queue Size: 0

Interface Name: PVL Mount Notification Interface

Interface ID: 007ff347-e533-1cc6-b22d-02608c2cedf4

Interface Version: 1

Authentication Mechanisms

☒ KRB5 ☒ UNIX

Interface Name: Realtime Monitor Interface

Interface ID: 80c9a256-2f13-11d3-a0c8-000001341966

Interface Version: 1

Authentication Mechanisms

☒ KRB5 ☒ UNIX

Interface Name: Client Interface

Interface ID: 32ba9692-4667-11d6-aa3a-0004ac49692b

Interface Version: 2

Authentication Mechanisms

☒ KRB5 ☒ UNIX

Interface Name: Account Validation Interface

Interface ID: 647f22a8-a1e9-11d3-a739-000001341966

Interface Version: 1

Authentication Mechanisms

☒ KRB5 ☒ UNIX

Update Delete Start Over Clone (partial) Dismiss

The **Interface Controls** section of the *Server Configuration* window is common to all servers. In the example window above, the server displayed is a Core Server.

Field descriptions.

Maximum Connections

The maximum number of clients that this server can service at one time. This value should be set based on the anticipated number of concurrent clients. Too large a value may cause unwarranted memory usage. Too small a value will mean that some clients may not be able to connect. This value should be chosen such that it is equal to between one and two times the number of clients that can be initiating connections simultaneously.

Thread Pool Size

The number of threads this server spawns to handle client requests. If necessary, the default values can be changed when defining servers. Too large a value may consume server memory for no purpose. Too small a value could mean that some client requests are rejected or don't get serviced in a timely manner. The Thread Pool Size should be equal to or larger than the number of expected simultaneous requests.

Request Queue Size

The maximum number of requests that can be queued waiting for request threads. If the workload increases so that this value is exceeded, requests will be rejected by the server rather than queued for processing. A value of zero means to use the default queue size of 20.

Note: See *Section 3.1.2, "Tuning the System Manager RPC thread pool and request queue sizes"* for information on tuning the RPC thread pool and request queue sizes for the System Manager.

The following three interfaces fields provide the information required for the server to communicate with other HPSS servers and clients over the HPSS RPC network interface. Each server type is configured with one or more interfaces. With the exception of the Authentication Mechanisms, the administrator cannot override the default values for each interface when creating or modifying a server configuration. Each interface consists of the following fields:

Interface Name

The descriptive name of the interface.

Interface Id

The UUID that identifies this interface.

Interface Version

The interface version number.

Authentication Mechanisms

The authentication mechanisms from which the interface will accept credentials. One or both of the mechanisms can be checked (at least one should be checked or that interface will be unusable). Each interface supports the following authentication mechanisms:

- KRB5 indicates that the interface will support Kerberos 5 authentication.
- UNIX indicates that the interface will support UNIX authentication.

5.2.1.4. Security controls

The security controls fields define the settings for authenticated communication between the server and other HPSS servers and clients.

The screenshot shows the 'Core Server Configuration' window with the 'Security Controls' tab selected. The window has a menu bar with 'File', 'Edit', and 'Help'. The 'Basic Controls' section includes fields for 'Server Name' (Core Server (Subsystem 1)), 'Server ID' (7), 'UUID' (4ca0aad4-8e19-11e5-83b3-005056b7004f), 'Server Type' (CORE), 'Storage Subsystem' (Subsystem #1), and 'Log Policy' (Core Log Policy). The 'Security Controls' section includes 'Principal Name' (hpsscore), 'Protection Level' (Connect), and two 'Authentication Service Configuration' sections. The first section has 'Mechanism' (UNIX), 'Authenticator Type' (Keytab), and 'Authenticator' (auth_keytab:/var/hpss/etc/hpss.unix.keytab). The second section has 'Mechanism' (KRB5), 'Authenticator Type' (Keytab), and 'Authenticator' (auth_keytab:/var/hpss/etc/hpss.keytab). At the bottom are buttons for 'Update', 'Delete', 'Start Over', 'Clone (partial)', and 'Dismiss'.

The **Security Controls** section of the *Server Configuration* window is common to all servers. In the example window above, the server displayed is a Core Server.

Field descriptions.

Principal Name

The name of the principal the server will use to authenticate.

Protection Level

The level of protection that will be provided for communication with peer applications. The higher the level of protection, the more encryption and overhead required in communications with peers. The levels, from lowest to highest, are as follows:

Connect

Performs authentication only when the client establishes a connection with the server.

Packet

Ensures that all data received is from the expected client.

Packet Integrity

Verifies that none of the data transferred between client and server has been modified.

Packet Privacy

Verifies that none of the data transferred between client and server has been modified and also encrypts the data transferred between client and server.

Authentication Service Configuration

Each server can support up to two Authentication Services. The following fields are used to define each authentication service configured for a server.

Mechanism

The authentication mechanism to use when passing identity information in communications to HPSS components.

- `KRB5` indicates that the server will use Kerberos 5 authentication.
- `UNIX` indicates that the server will use UNIX authentication.
- `Not Configured` indicates that an authentication service has not been configured for this slot. At least one of the authentication service slots must be configured.

Authenticator Type

The type of authenticator specified in the **Authenticator** field. The types are:

- `Not Configured` indicates that an authenticator has not been configured for this slot. If a mechanism is specified, an authenticator type must also be specified.
- `None` indicates no authenticator is supplied for this mechanism. This is appropriate for UNIX authentication if no keytab is used. The server's credentials will be its current UNIX identity.
- `Keytab` indicates that the authenticator is the path to a keytab file. For Kerberos authentication this is a keytab file created with Kerberos utilities. For UNIX authentication this is a keytab file created with the `hpss_unix_keytab` utility. See its man page for details. Each server can have its own keytab file, or all the servers can share a single keytab file. It is recommended that one keytab file be used for all of the servers on any given host.

Authenticator

The argument passed to the authentication mechanism indicated by the **Authenticator Type** configuration variable and used to validate communications. If it is a keytab, the server must

have read access to the keytab file. Other access permissions should not be set on this file or security can be breached. For the Not Configured or None values of the **Authenticator Type**, this field can be left blank.

5.2.1.5. Audit policy

HPSS security provides a log for certain security-related events such as a change in an object's permissions, a change in an object's owner, a change of an object's name, a failed attempt to access an object, and several others (see *Field descriptions* below for a complete list). A server's Audit Policy controls which of these events trigger log messages. It can be configured from the **Audit Policy** tab of the *Server Configuration* window. It is possible to request that only failures or both successes and failures be logged.

For each audit event type, selecting **Failures** causes only failures to be logged, selecting **Total** logs every audit event of that type, and selecting neither causes nothing to be logged. Selecting both has the same effect as selecting **Total**.

Currently the only server which generates security object event records other than AUTH is the Core Server.

For more information on auditing, see the *Security audit* subsection of the *HPSS Installation Guide* (in the *HPSS planning* chapter).

Core Server Configuration

File Edit Help

Basic Controls

Server Name: Core Server (Subsystem 1)

Server ID: 7

UUID: 4ca0aad4-8e19-11e5-83b3-005056b7004f

Server Type: CORE

Storage Subsystem: Subsystem #1

Log Policy: Core Log Policy

Execution Controls **Interface Controls** **Security Controls** **Audit Policy** **Specific**

	Total	Failures		Total	Failures
AUTH	<input type="checkbox"/>	<input checked="" type="checkbox"/>	RENAME	<input type="checkbox"/>	<input checked="" type="checkbox"/>
CHMOD	<input type="checkbox"/>	<input checked="" type="checkbox"/>	RMDIR	<input type="checkbox"/>	<input checked="" type="checkbox"/>
CHOWN	<input type="checkbox"/>	<input checked="" type="checkbox"/>	UNLINK	<input type="checkbox"/>	<input checked="" type="checkbox"/>
CREATE	<input type="checkbox"/>	<input checked="" type="checkbox"/>	UTIME	<input type="checkbox"/>	<input checked="" type="checkbox"/>
LINK	<input type="checkbox"/>	<input checked="" type="checkbox"/>	ACL_SET	<input type="checkbox"/>	<input checked="" type="checkbox"/>
MKDIR	<input type="checkbox"/>	<input type="checkbox"/>	CHBFID	<input type="checkbox"/>	<input checked="" type="checkbox"/>
OPEN	<input type="checkbox"/>	<input checked="" type="checkbox"/>	BFSETATTRS	<input type="checkbox"/>	<input checked="" type="checkbox"/>

Update Delete Start Over Clone (partial) Dismiss

The **Audit Policy** section of the *Server Configuration* window is common to all servers. In the example window above, the server displayed is a Core Server.

Field descriptions.

AUTH

Authentication events.

CHMOD

Core Server object permission events.

CHOWN

Core Server object owner events.

CREATE

Core Server creation events.

LINK

Core Server hard link creation events.

MKDIR

Core Server directory creation events.

OPEN

Core Server bitfile open events.

RENAME

Core Server object rename events.

RMDIR

Core Server directory remove events.

UNLINK

Core Server object delete events.

UTIME

Core Server bitfile time modified events.

ACL_SET

Core Server access control list modification events.

CHBFID

Core Server change bitfile identifier events.

BFSETATTRS

Core Server set bitfile attribute events.

5.2.2. Core Server-specific configuration

A separate Core Server must be configured for each storage subsystem.

The screenshot shows the 'Core Server Configuration' window with the 'Specific' tab selected. The window has a menu bar with 'File', 'Edit', and 'Help'. The 'Basic Controls' section includes fields for 'Server Name' (Core Server (Subsystem 1)), 'Server ID' (7), 'UUID' (4ca0aad4-8e19-11e5-83b3-005056b7004f), 'Server Type' (CORE), 'Storage Subsystem' (Subsystem #1), and 'Log Policy' (Core Log Policy). The 'Execution Controls' section includes fields for 'Root Fileset Name' (FilesetRoot.2767), 'Root Fileset ID' (3272905551, 2960343120), 'Maximum Open Bitfiles' (2000), 'Maximum Active I/O Requests' (190), 'Maximum Active Copy Requests' (95), 'COS Change Retry Limit' (0), 'Tape Dismount Delay (seconds)' (30), 'Tape Handoff Delay (seconds)' (120), 'Tape Queue Method' (Offset), 'Tape Queue Active Set Size' (1024), 'Tape Max Reader Wait (seconds)' (300), 'PVL Max Connection Wait (seconds)' (300), and 'Fragment Smallest Block' (8 (2^3)). There are checkboxes for 'COS Copy to Disk', 'Log ENOSPACE Details (debug only)', and 'I/O Operations to Log (debug only)'. The 'I/O Operations to Log' section has a table with checkboxes for 'Read', 'Write', and 'Copy' for 'Disk' and 'Tape'. The 'Classes of IODs/IORs to Log (debug only)' section has checkboxes for 'SS Client IODs/IORs', 'Mover IODs/IORs', 'Rait Engine IODs/IORs (tape only)', 'Passive Replies', 'Mover Abort IODs', and 'Mover IOCTLS'. At the bottom, there are buttons for 'Update', 'Delete', 'Start Over', 'Clone (partial)', and 'Dismiss'.

The **Specific** tab of the *Core Server Configuration* window allows you to view and update the type-specific configuration for the Core Server.

The settings on this screen can be changed while the Core Server is running by changing the value on the screen, clicking the **Update** button to update the metadata, and then clicking the **Apply Config** button on the *HPSS Health and Status* screen to reinitialize the appropriate Core Server. The Core Server rereads the metadata and changes its internal settings. See Section 5.8.2, “Reinitializing a server” for more information.



Changes to switches under **I/O Operations to Log (debug only)** and **Classes of IODs/ IORs to Log (debug only)** take effect immediately after clicking the **Update** button and do not require server reinitialization.

Field descriptions.

Root Fileset Name

Core Servers can create and support multiple filesets, but an initial fileset is required for each Core Server. The **Root Fileset Name** designates which of these filesets will be used by the Core Server to resolve the pathname "/" (forward slash) in the subsystem. Other filesets served by this Core Server, or from other Core Servers in the HPSS system may be joined to this fileset by junctions.

Root Fileset ID

The Fileset ID of the fileset named in **Root Fileset Name**.

Maximum Open Bitfiles

The maximum number of bitfiles that can be open simultaneously. This value is augmented by the number of open files required for tape aggregation. See *Section 5.2.2.3, "Core Server max open files configuration"* for more information.

Maximum Active I/O Requests

The maximum number of simultaneous I/O requests allowed. (See *Section 5.2.2.2, "Core Server I/O limit configuration"*. for a detailed description of how this field interacts with other configuration fields)

Maximum Active Copy Requests

The maximum number of simultaneous copy requests allowed. (See *Section 5.2.2.2, "Core Server I/O limit configuration"*. for a detailed description of how this field interacts with other configuration fields)

COS Change Retry Limit

This is the maximum number of attempts that will be made to change the Class of Service of a file. If set to "0", COS change will continue to be attempted until it is successful. If a positive value is provided, the COS change request will be dropped after it has failed the configured number of times.

Tape Dismount Delay (seconds)

The amount of time, in seconds, a mounted tape volume will remain idle before being dismounted by the Core Server. Larger values may reduce undesirable tape dismount/remount events, at the expense of lower tape drive utilization.

Tape Handoff Delay (seconds)

The amount of time, in seconds, a mounted tape will be held in a client's session before become eligible to be handed off to another client that wishes to use the tape.

Tape Queue Method

This field may be set to either `FIFO` or `Offset`. If set to `FIFO`, the Core Server will process tape requests per VV in a FIFO fashion without doing any intelligent ordering. This somewhat mimics legacy behavior and can be useful for sites which want to closely manage their own ordering. If set to `Offset`, the Core Server will attempt to reorganize the VV session queue in order to improve performance. If available, this may use drive technologies such as Recommended Access Ordering (RAO).

If a file called `nodrive` exists in the HPSS tmp folder (`HPSS_PATH_TMP`), its presence will disable RAO-based sorting of tape requests and instead use strict offset ordering. If a file called `noschedule` exists in this folder it will disable all scheduling, resulting in a FIFO ordering for tape requests from that point until the file is removed. Note that this will not impact any existing schedules.

Individual VVs may also be toggled to FIFO scheduling using the **FIFO** check box on the VV information window. See the *Section 8.5.4.2, “Core Server Tape Volume information window”* for more information.

Tape Queue Active Set Size

Number of sessions which may be scheduled together. This yields a compromise between scheduling efficiency and potential wait times since this size is the maximum number of items within that set that may be scheduled ahead of any other request in that set. This may also be limited by the hardware limitations of Recommended Access Ordering (RAO) if the drive a tape is mounted on supports RAO.

Tape Max Reader Wait (seconds)

This is the amount of time, in seconds, that readers of a tape may spend waiting on writers before the scheduling priority shifts to prefer readers over writers for some period of time. There is a minimum of five minutes between reader/writer priority changes. Setting this to zero disables preference switching and will cause writers to always be preferred.

PVL Max Connection Wait (seconds)

The amount of time, in seconds, the Core Server will wait to connect to the PVL before declaring an error in pending PVL jobs.

Fragment Smallest Block

Fragment Smallest Block sets the boundary used to determine where to remove excess clusters from allocations that have excess space. The smallest block returned to the free space map from the excess space at the end of an extent will be this size or larger.

COS Copy to disk

If ON, all copy operations associated with COS changes will be directed to disk if the hierarchy has a disk storage class at its top level, and will subsequently be migrated down the hierarchy. If OFF, COS changes are not automatically copied to the disk level. They will instead go to the top-level tape storage class and then migrate down from there.

Utilizing COS Copy to disk allows migration to occur from that disk COS according to the disk's migration policy with features such as ordered migration and tape aggregation. Otherwise, files are written individually to tape for each COS stream. A large number of COS change streams could cause data to be spread out over a large number of tape volumes. Having this feature **off** is generally inefficient for small files or a large number of COS change streams, and could result in poor COS change performance or recall.

COS Copy to disk is recommended for systems that contain significant numbers of small files.

Log ENOSPACE Details (debug only)

If ON, the Core Server logs detailed debug class log messages about disk and tape allocation requests that end in `HPSS_ENOSPACE`. In normal operation, these details are too voluminous to

log economically, so they are not logged by default. But if the cause of an ENOSPACE condition is not evident, setting this switch will enable logging of all of the details. Because of the volume of information logged, leaving this switch on for long periods of time is not recommended.

I/O Operations to Log (debug only), Classes of IODs/IORs to Log (debug only)

These switches control optional logging in the Core Server. The Core Server always logs IODs and IORs when I/O operations fail, but does not log IODs and IORs for successful operations. These switches allow the administrator to turn on logging of IODs and IORs in successful tape and disk I/O operations for diagnostic purposes.

Normally these switches should be off to avoid logging large amounts of information that won't be useful. Changes to these switch settings in the managed object take effect immediately.

Reinitialization of the server is not necessary.

5.2.2.1. Additional Core Server configuration

Certain aspects of the operation of the Core Server can be controlled by setting environment variables. The server uses built-in default values for these settings, but if the environment variables can be found in the server's environment, the server uses those values. The following is a list of the variables and the aspects of the server's operation they control. Since the environment is common to all subsystems, all Core Servers in an HPSS installation are subject to these values.

HPSS_MAX_NUM_OPEN_FILES_FOR_AGGR

Adjusts the maximum number of files that may be open by the Core Server when it is assembling migration file aggregates. If this variable is not present in the environment, the server's default value is 20000. See *Section 5.2.2.3, "Core Server max open files configuration"* for more information.

HPSS_SPEC_USER_LIST_LENGTH

Adjusts the length of the Core Server's internal list of users who have delete only permission in AUTHZACL table. If this variable is not present in the environment, the server's default value is 8.

HPSS_CORE_ETIME_POLICY

Selects the algorithm the Core Server will use whenever it computes the etime.

There are three possible algorithms: CORE_ETIME_POLICY_NOETIME, CORE_ETIME_POLICY_RELATIVE, and CORE_ETIME_POLICY_STRICTETIME. CORE_ETIME_POLICY_NOETIME indicates that the etime is not to be calculated or set. The CORE_ETIME_POLICY_RELATIVE value indicates that the Unix Relative etime algorithm is to be used when computing the etime. The CORE_ETIME_POLICY_STRICTETIME value indicates that the Core Server is to use the traditional Unix algorithm when computing the etime. Note that the traditional Unix algorithm (CORE_ETIME_POLICY_STRICTETIME) can cause performance issues and these performance issues are the reason that the Relative etime algorithm was developed. Further note that the CORE_ETIME_POLICY_RELATIVE is recommended and is the default algorithm used by the Core Server when computing the etime.

HPSS_CORE_REPACK_OUTPUT

Enables or disables the repack output segregation feature. If this environment variable is not present in the environment, or has the value "on", files moved from one tape virtual volume to another by **repack** will be written to tapes that contain only files written by the **repack**. If this

environment variable has the value "off", files moved from one tape virtual volume to another by **repack** may be mixed on output tapes with files written by the migration subsystem, or directly to tape by users. When set to "on", this feature can be ignored on a per-**repack** basis with the `--no-dest-tape-segregation` command line option. See the **repack** man page for details.

HPSS_CORE_SEG_CACHE_IDLE_TIME

Adjusts the time, in seconds, a disk or tape storage segment can remain idle in the Core Server's in-memory cache before being purged. If this variable is not present in the environment, the server's default value is 60 seconds.

HPSS_CORE_TAPE_CACHE_IDLE_TIME

Adjusts the time, in seconds, a tape's in-memory cache entries can remain idle before being purged. This value controls the tape storage map, VV and PV caches. If this variable is not present in the environment, the server's default value is 300.

HPSS_CORE_LARGE_SEG_THRESHOLD

Controls the way tape storage space is allocated when large tape storage segments are written. The value of this environment variable is a time, measured in seconds. If the segment to be written would take longer to write than this time threshold, then the segment is written to a tape VV that has sufficient free space. This may mean the system will mount a different tape if it had been writing a tape that is nearly full. If the segment will take less time to write than this threshold, then the segment is written to any available tape, including a tape that will soon reach End Of Media (EOM). The server's default value is 30 seconds. The server uses the Transfer Rate information from the storage class when estimating the time required to write a segment.

This threshold value also applies when **repack** is moving tape storage segments from one tape to another. Large tape storage segments will be moved to a tape that has sufficient free space if the time to move the segment exceeds the threshold value. Tapes selected for repack output are also subject to the `HPSS_CORE_REPACK_OUTPUT` limitation. When **repack** creates new tape aggregates using small files, it writes to any available tape, subject to the rules defined by `HPSS_CORE_REPACK_OUTPUT`.

HPSS SOCK_KEEPIDLE_SECS

Adjusts the time, in seconds, that a socket connected to a Mover will be idle before TCP pings the Mover. If the pings are not answered, TCP returns an error to the Core Server, after about 11 more minutes. The error allows the Core Server to recover control of a disk or tape transfer if the Mover crashes or if the node the Mover runs on crashes. Note that this timer does not help the Core Server regain control of a transfer if the Mover hangs up on a storage device. The default value of this setting is 7200 seconds (2 hours).

HPSS_CORE_DISKSEGCACHE_MEM_PERCENT

The Core Server's disk segment cache occupies, by default, 1% of total system memory. For example, on an HPSS Core Server machine with 256 GB of memory, the disk segment cache will use up to 2.56 GB of memory. Increasing this value beyond 1% may lead to increased file creation and deletion rates, at the cost of increased memory usage. Decreasing this value below 1% may lead to decreased file creation and deletion rates, with the benefit of decreased memory usage.



Advice

For more information on tape ordered recall, see Section 13.1, "Overview of TOR".

5.2.2.2. Core Server I/O limit configuration

The number of concurrent IO operations allowed in each Core Server is influenced by several factors, including three parameters from the Core Server configuration. Two of these are in the **Core Server-specific configuration** tab of the *Core Server Configuration* window in SSM:

- Maximum Active I/O Reqs
- Maximum Active Copy Reqs

The other is in the **Interface Controls** tab:

- Thread Pool Size

The Core Server tracks the current number of IO requests in three categories:

- Active Total IO
- Active Copy IO
- Migrate IO

Active Copy IO includes stages, internal copies of one hpss file to another, and COS changes. Active Total IO includes all of the Active Copy IO plus reads by client apps from hpss and writes to hpss by client apps. Migrate IO includes only migrations.

Active Copy IO is a subset of Active Total IO. Migrate IO is in a class by itself; it is neither a subset nor a superset of Total IO nor of Copy IO.

The maximum values allowed for Active Total IO and Active Copy IO may be configured by changing the "Maximum Active I/O Requests" and "Maximum Active Copy Requests" parameters in the **Core Server-specific configuration**, respectively. There is no configurable maximum for Migrate IO.

IO requests submitted after these maximum values are reached will be rejected. However, up to 2000 additional stage requests can also be submitted to be run as background stages, regardless of the configured max values. These background stage requests are queued until there is room to run them, that is, until one or more of them can be handled without going over our maximum Active Copy IO limit. These background stage requests are not counted in the Active Copy IO until they are released from the queue and actually submitted for the stage.

The Core Server may adjust the actual values it uses for the Maximum Active Total IO and Maximum Active Copy IO based on the "Thread Pool Size" parameter in the *Core Server Configuration*. The Thread Pool Size specifies the total number of threads the Core Server may create. It must be large enough to allow for reasonable operation, and it must be large enough for the Maximum Active Total IO plus 20 threads reserved for Core Server internal use.

The bare minimum for "reasonable operation" is 25 threads. If Thread Pool Size is less than 25, the Core Server changes its actual Max Active Copy IO and Max Active Total IO to very small values, no greater than 10 each, and in some cases to 0 each.

If Thread Pool Size is greater than 25, but is still not large enough for the Maximum Active Total IO plus 20 threads reserved for Core Server internal use, then the Core Server reduces the actual max values it will use in this way:

1. It sets its actual Maximum Active Total IO to 20 less than the Thread Pool Size:
 - Actual Max Active Total IO = Thread Pool Size - 20
2. It calculates the percentage of Active Copy IO with regard to Active Total IO which was requested in the configuration:

```
percentage = (configured Max Active Copy IO) /  
              (configured Max Active Total IO)
```

3. It sets its actual Maximum Active Copy IO to this percentage of its actual Maximum Active Total IO:

```
Actual Max Active Copy IO = percentage * Actual Max Active Total IO
```

So the site should configure Max Active Copy IO large enough to handle its expected peak concurrent number of stages, internal file copies, and COS changes.

It should configure Max Active Total IO large enough to handle all of the Active Copy IO plus enough for all the expected concurrent reads and writes.

It should configure Thread Pool Size large enough to handle Max Active Total IO, plus the number of migrations it expects to have, plus at least 20 for Core Server internal threads.

Changes to the Thread Pool Size take effect only after clicking **Apply Config** on the *HPSS Health and Status* screen or on a Core Server restart or reinitialization.

The SSM **Core Server-specific configuration** window shows the values of the Max Active Total IO and Max Active Copy IO as they existed in the **Core Server-specific configuration** based on the last time a configuration update was applied.

To see the adjusted values, look in the *Core Server Configuration* window or the **Core Server Information** window in SSM. Alternatively, use the scrub limits command, or issue a SIGHUP to the Core Server.

Here is an example of the scrub limits command:

```
scrub> limits  
      Max Open Bitfiles =          14000  
      Cur Open Bitfiles =           12  
      Max Active IO =             600  
      Cur Active IO =              0  
      Max Active Copy IO =         600  
      Cur Active Copy IO =          0  
      Max Connections =          1600  
      Cur Connections =           3  
      Tape Queue Method =         Offset
```

If a SIGHUP is sent to the Core Server, examine the "BFS IO State" section of the resulting dump file. Here is an example of this section:

```

===== BFS IO State =====
Total Number of Active IO Requests      210
Max Number of IO Requests Allowed       300
Total Number of Active Copy Requests    180
Max Number of Copy Requests Allowed     200
Total Number of Active Migrate Requests  40

```

Applications can retrieve the values for the configured limits via the API using the `hpss_GetSubSysLimits` function.

5.2.2.3. Core Server max open files configuration

There are two settings used in determining the maximum number of concurrently open bitfiles allowed in HPSS: one setting for files opened by tape aggregation and one setting for files opened for any other purpose.

The values of the two settings are added together to reach the total number of files allowed to be open at any one time. In enforcing this limit, the Core Server does not check whether a file is being opened for tape aggregation or for some other reason. Having the two separate values is only a coarse means of helping the administrator estimate the needs of the system based on the system configuration and workload.

The limit for the max number of files opened for purposes other than tape aggregation is set in the **Core Server-specific configuration** in the "Maximum Open Bitfiles" field.

The limit for the max number of files opened concurrently for tape aggregation is hard-coded into the Core Server as 20,000; however, it may be changed by setting the environment variable `HPSS_MAX_NUM_OPEN_FILES_FOR_AGGR` to the desired value. The Core Server may use a value less than this limit based on its calculation of what is actually needed for aggregation.

On startup, the Core Server computes how many open files would be required for tape aggregation if an aggregate migration were executing concurrently on all possible combinations of storage classes and hierarchies. This depends upon several factors:

1. The number of disk storage classes.
2. The number of hierarchies that use each disk storage class.
3. The copy count in each affected hierarchy.
4. The "Total Migration Streams" defined in the migration policy of each storage class.
5. The number of batch migration requests allowed concurrently per stream, which is 2.
6. The max number of files allowed per aggregate, which is 50,000.

The calculation for the number of open files required for aggregate migration for one storage class as used in one hierarchy would be:

```
50,000 * stream count * 2 * copy count
```

The Core Server adds up these values for all disk storage classes as they are used in all hierarchies. The resulting value is not allowed to exceed `HPSS_MAX_NUM_OPEN_FILES_FOR_AGGR`, or,

if the `HPSS_MAX_NUM_OPEN_FILES_FOR_AGGR` variable is not set, then it is not allowed to exceed 20,000.

This value is then added to the value of "Maximum Open Bitfiles" from the **Core Server-specific configuration** for the actual maximum open files allowed for the entire system.

To be recognized, changes in either value require a reinitialization or, sometimes, a restart of the Core Server. If the new total value would be less than the allocated size of the open context list, it can be implemented with a reinitialization. Changes to a larger value require a Core Server restart. The allocated size of the open context list is 4 times the initial value of the maximum number of open files.

A SIGHUP dump of the Core Server shows the current maximum open files allowed and the actual size of the open context list:

```
Maximum open files      = 6000
Allocated Size of Open Context List  = 24000
```

When the Core Server starts up, it issues a log message saying how many additional files it allowed for aggregation. The message looks something like this:

```
Tape aggregation on, adjusting open file cache: additional open files added=20000
```

If at startup the Core Server realizes it needs more open files for aggregation than allowed by the value of `HPSS_MAX_NUM_OPEN_FILES_FOR_AGGR` (or more than 20,000, if `HPSS_MAX_NUM_OPEN_FILES_FOR_AGGR` is not set), then it issues a warning message that looks something like this:

```
Tape aggregation on, adjusting open file cache:
additional open files needed=5700000,
limited by max = 20000 (Too many open bitfiles)
```

5.2.3. Gatekeeper-specific configuration

To use a Gatekeeper for gatekeeping services the Gatekeeper must be configured into one or more storage subsystems (see *Section 4.2.3, "Storage Subsystem Configuration window"*). To associate a Gatekeeper with a storage subsystem, the Gatekeeper must be selected on the *Storage Subsystem Configuration* window.

To use the Gatekeeper for account validation services, account validation must be enabled in the accounting policy (see *Section 14.2, "Accounting"*). If account validation is configured to be ON, then at least one Gatekeeper must be configured even if the site does not wish to do gatekeeping. In this case, the Gatekeeper will only be used for validating accounts. Otherwise a Gatekeeper may be used for both gatekeeping and account validation. If multiple Gatekeepers are configured, then any Gatekeeper may be contacted for account validation requests.

Note: If a Gatekeeper is configured, then it will either need to be running or marked non-executable for HPSS Client API requests to succeed in the Core Server (even if no gatekeeping nor account validation is occurring); this is due to the HPSS Client API performing internal accounting initialization.

For a detailed description of the Gatekeeper, refer to the *Gatekeeper* section of the *HPSS Installation Guide* (in Chapter *HPSS planning*).

The screenshot shows the 'Gatekeeper Configuration' window. The 'Basic Controls' tab is active, displaying fields for Server Name (Gatekeeper), Server ID (5), UUID (095067d2-8e01-11e5-83b3-005056b7004f), Server Type (GK), and Log Policy (Default). Below this, the 'Execution Controls' tab is selected, showing 'Default Wait Time (seconds)' (10) and 'Site Policy Pathname (UNIX)' (/var/hpss/gk/gksitepolicy). At the bottom, there are buttons for 'Update', 'Delete', 'Start Over', 'Clone (partial)', and 'Dismiss'.

Field descriptions.

Default Wait Time (seconds)

The number of seconds the client must wait before retrying a request. This value must be greater than zero and is used if the Gatekeeping Site Interface returns a wait time of zero for the create, open, or stage request being retried.

Site Policy Pathname (UNIX)

The contents of this file will be defined by the site and should be coordinated with the Gatekeeping Site Interface. For example, a site may want to define the types of requests being monitored, the maximum number of opens per user, or a path name to a log file. The **Site Policy Pathname** will be passed to the `gk_site_Init` routine upon Gatekeeper startup. If a site wishes to make use of the **Site Policy Pathname** file, then the site will need to add code to the Gatekeeping Site Interface library to read the **Site Policy Pathname** file and interpret the data accordingly. Refer to the *Site Interfaces* chapter of the *HPSS Programmer's Reference* for details on the site's interface library.

5.2.4. Location Server additional configuration

The Location Server does not have a specific server configuration in the *Server Configuration* window. However, some additional configuration outside SSM is required whenever a Location Server is added or modified. Each client application which accesses HPSS must first contact the Location Server using the Location Server's RPC endpoints, which must be listed in the `/var/hpss/etc/ep.conf` file to be available to client applications. The **hpss_bld_ep** utility reads the Location

Server configuration and creates the `ep.conf` file. Refer to *The /var/hpss files* appendix of the *HPSS Installation Guide* for more information on the `ep.conf` file.

The **hpss_bld_ep** utility must only run on the HPSS root subsystem machine. To invoke the utility:

```
% /opt/hpss/config/hpss_bld_ep -v -f /var/hpss/etc/ep.conf -r <UNIX/LDAP
Realm Name> -c <schema name> -d <DB2 configuration database> add
```

For example:

```
% /opt/hpss/config/hpss_bld_ep -v -f /var/hpss/etc/ep.conf -r
hpss.acme.com -c hpss -d cfg add
```

The location server client-side library will cache location information locally on client systems. The cache is located in the HPSS application's home directory (`~/.hpss_ls_cache`) with permissions 0600 and is used to improve the startup time of client applications. This feature can be disabled by setting `LS_DISABLE_LOCAL_CACHE=1` in the client system's HPSS `env.conf`.

5.2.5. Migration/Purge Server (MPS)-specific configuration

The screenshot shows a window titled "Migration/Purge Server Configuration" with a menu bar (File, Edit, Help). The window is divided into two main sections. The top section, labeled "Basic Controls", contains fields for Server Name (Migration/Purge Server), Server ID (9), UUID (a333e356-8e1a-11e5-83b3-005056b7004f), Server Type (MPS), Storage Subsystem (Subsystem #1), and Log Policy (Default). The bottom section, labeled "Specific", contains fields for Storage Class Update Interval (60 seconds), Core Server API Failures Allowed (3), and MPS Report File Name (path/prefix). At the bottom of the window are buttons for Update, Delete, Start Over, Clone (partial), and Dismiss.

Basic Controls
Server Name: Migration/Purge Server
Server ID: 9
UUID: a333e356-8e1a-11e5-83b3-005056b7004f
Server Type: MPS
Storage Subsystem: Subsystem #1
Log Policy: Default

Specific
Storage Class Update Interval: 60 seconds
Core Server API Failures Allowed: 3
MPS Report File Name (path/prefix):

Buttons: Update, Delete, Start Over, Clone (partial), Dismiss

This tab of the *Migration/Purge Server Configuration* window allows you to update and view the type-specific configuration parameters for a Migration/Purge Server (MPS).



*The server, if already running, must be reinitialized or recycled in order for the changes to become effective. It is recommended to use the **Apply Config** button on the HPSS Health and Status window to reinitialize the server.*

Field descriptions.

Storage Class Update Interval

The interval that indicates how often the MPS will query the Core Server in its subsystem to get the latest storage class statistics. This is also the interval the MPS uses to check whether it needs to initiate a purge operation on the storage class based on the associated purge policy. The valid range for this field is 10 - 600 seconds.

Core Server API Failures Allowed

The maximum number of consecutive Core Server API failures allowed before MPS aborts a migration or purge run. This applies to both disk and tape migration as well as disk purge. If this limit is reached during a disk migration run, migration skips to the next hierarchy to which that disk storage class belongs. Once the run has attempted to migrate all such hierarchies the run will end. If this limit is reached during a disk purge or tape migration run, the migration or purge run aborts.

MPS Report File Name (path/prefix)

A prefix string used by the MPS to construct a migration report file name. The full file name will consist of this string with a date string and subsystem ID appended to it. If this field is left blank, no migration reports will be generated. If a usable path prefix is not specified, the location of the migration report files may be unpredictable. A new MPS report file is started daily at midnight local time.

5.2.6. Mover-specific configuration

5.2.6.1. Mover-specific configuration window

The screenshot shows the 'Mover Configuration' window with the 'Specific' tab selected. The window has a title bar with standard OS controls and a menu bar with 'File', 'Edit', and 'Help'. The 'Basic Controls' section contains fields for 'Server Name' (Mover (hpss-dev-dus03)), 'Server ID' (10), 'UUID' (d14e0ad2-8e1a-11e5-83b3-005056b7004f), 'Server Type' (MOVER), and 'Log Policy' (Default). Below this is a tabbed interface with 'Execution Controls', 'Interface Controls', 'Security Controls', 'Audit Policy', and 'Specific'. The 'Specific' tab contains fields for 'Encryption Key' (0xb60a080d5075b5b4) with a 'Generate' button, 'Buffer Size' (1048576), 'TCP Path Name' (/opt/hpss/bin/hpss_mvr_tcp), 'Hostname' (hpss-dev-dus03.ccs.ornl.gov), 'TCP Port' (5001), 'Port Range Start' (0), and 'Port Range End' (0). A 'Data Hostnames' section has a list of 8 input fields, with the first one containing 'hpss-dev-dus03.ccs.ornl.gov'. At the bottom are buttons for 'Update', 'Delete', 'Start Over', 'Clone (partial)', and 'Dismiss'.

Mover Configuration

File Edit Help

Basic Controls

Server Name: Mover (hpss-dev-dus03)

Server ID: 10

UUID: d14e0ad2-8e1a-11e5-83b3-005056b7004f

Server Type: MOVER

Log Policy: Default

Execution Controls Interface Controls Security Controls Audit Policy **Specific**

Encryption Key: 0xb60a080d5075b5b4 [Generate](#)

Buffer Size: 1048576

TCP Path Name: /opt/hpss/bin/hpss_mvr_tcp

Hostname: hpss-dev-dus03.ccs.ornl.gov

TCP Port: 5001

Port Range Start: 0

Port Range End: 0

Data Hostnames

0: hpss-dev-dus03.ccs.ornl.gov

1:

2:

3:

4:

5:

6:

7:

[Update](#) [Delete](#) [Start Over](#) [Clone \(partial\)](#) [Dismiss](#)

This tab of the *Mover Configuration* window allows you to update and view the type-specific configuration for the Mover.

The server must be recycled in order for changes to become effective.

The Mover cannot be configured entirely from the SSM window. See *Section 5.2.6.2, “Additional Mover configuration”* for further necessary configuration.

Field descriptions.

Encryption Key

An encryption key used to secure the Mover’s interface with the Core Server. A specific value may be entered in the field or a random value may be generated by clicking on the **Generate** button. The default value is "0". A nonzero value must be configured to allow client access to the Mover’s data interface. See *Section 5.2.6.2, “Additional Mover configuration”*, for further information about the encryption key.

Buffer Size

The buffer size used for double buffering during data transfers. The default buffer size is set to 1,048,576 bytes (1 MB). This value should be tuned based on device and networking configuration and usage. The trade-off for this value is that large buffer sizes will use more system memory and may be inefficient for small transfers. For example, if the Mover buffer size is 4 MB, but client requests are 512 KB, the Mover will not achieve any double buffering benefit because the entire amount of the transfer fits in one Mover buffer. A smaller buffer size will cause device and network I/O to be interrupted more often, usually resulting in reduced throughput rates for all but the smallest transfers.

The minimum Mover buffer size is the smallest block size for any device the Mover will handle. The maximum value will be bounded by the available system memory and the number of concurrent Mover requests anticipated.

TCP Path Name

The pathname of the Mover TCP/IP listen executable. The default value is `/opt/hpss/bin/hpss_mvr_tcp`. The TCP Movers currently supported are listed below. All TCP Movers support common disk/tape interfaces, TCP/IP, and shared memory data transfers.

Table 5.1. Mover TCP pathname options

Name	Options supported
hpss_mvr_tcp	Standard disk/tape devices - tape devices that use the Native SCSI device driver

Hostname

The name of the host interface used for Mover control communication. This must be a valid hostname for one of the network interfaces on the Mover node. The **Execute Hostname** of the *Core Server Configuration* window is used as the default. If the Mover is running remotely, this field must correspond to a network interface on the remote node, whereas the **Execute Hostname** set in the Mover’s basic configuration corresponds to a network interface on which the administrative portion of the Mover runs.

Data Hostnames

The host network interface names to be used by the Mover when transferring data. There can be up to eight data hostnames specified. These must be valid hostnames for network interfaces on

the Mover node. The default value of the first hostname is the **Execute Hostname** of the *Core Server Configuration* window. This field must correspond to a network interface on the remote node. These network interface addresses are utilized during migration/stage to allow the Mover to distribute connections and data across all the available network interfaces.

It is important to keep in mind that these hostnames are used by the Mover when it must supply network endpoints to be used to initiate connections from other Movers. These addresses are not used in the setup of connections between client and Mover (with the exception of the PData Push Protocol).

TCP Port

The TCP/IP port number used by the administrative interface portion of the Mover for the TCP/IP listen process to receive connections. It should be a value over 5000. The default value is 5001. The port number must be unique for processes running on the same node.

The Mover will internally use the port one greater than the one configured in this field on the remote node during initialization (for example, if 5001 is entered, port 5002 is used on the remote node); therefore available port ranges on both nodes must be taken into consideration when selecting this value.

Port Range Start

The beginning of a range of local TCP/IP port numbers to be used by the Mover when connecting to clients. Valid values are zero or any valid TCP port number to which the Mover may bind (that is less than or equal to the value of **Port Range End**). The default value is "0". Use of particular port ranges may be required by some sites for communication across a firewall. If this value is set to zero, the operating system will select the port number; otherwise the Mover selects a local port number between **Port Range Start** and **Port Range End** (inclusive). If nonzero, this field must be less than or equal to **Port Range End**. If this field is zero, **Port Range End** field must also be zero.

Port Range End

Used in conjunction with **Port Range Start** (see above). Valid values are zero or any TCP port number to which the Mover may bind (that is greater than or equal to the value of **Port Range Start**). The default value is "0". If nonzero, this field must be equal or greater than **Port Range Start**. If this field is zero, **Port Range Start** field must also be zero.

5.2.6.2. Additional Mover configuration

The part of the Mover that handles accessing configuration metadata and servicing the Mover administrative interface runs on the Core Server platform, while the part of the Mover (with which the PVL and Core Servers communicate) that manages the storage devices and performs data transfers runs on a remote node. To support this mode of operation, there is some additional configuration which must be done on the remote node.

This additional configuration may be performed using **mkhpss**. See *the section called "Setting up remote Movers with mkhpss"*.

Alternately, the configuration may be performed manually using the instructions from the following sections:

- *the section called "/etc/services, /etc/inetd.conf, and /etc/xinetd.d"*.

- *the section called “The Mover encryption key files”.*
- *the section called “/var/hpss/etc files required for remote Mover”.*
- *the section called “System configuration parameters on Linux”.*

The Mover also responds to an environment variable, `MVR_CLIENT_TIMEOUT`. This environment variable may be set in `/var/hpss/etc/env.conf` on each remote Mover system. It represents the amount of time, in seconds, that the Mover will wait for messages from clients before timing out. The default timeout value of 60 minutes should be used in most circumstances. This should give plenty of time for an average tape drive to write the default maximum size aggregate.

If the max aggregate size is increased in the migration policy, this timeout may need to increase. The suggested value is about 133% of the time it takes to write that aggregate to tape. For example, a max aggregate size of 1TB with a transfer rate of 250MB/s should be set to 4000 (the duration of the transfer) * (4/3), for a resulting timeout of 5333 seconds.

Note that this timeout value affects Mover behavior in other ways, such as in error scenarios causing the Mover to require additional time to identify dropped clients and recover from it.

Similarly, the Mover responds to a `MVR_PASSIVE_TIMEOUT` environment variable. This environment variable may be set in `/var/hpss/etc/env.conf` and can be used to change the amount of time, in seconds, that a passive Mover process will wait before being contacted. This can be useful for stage operations which may have to wait a significant amount of time before being serviced. Consult with HPSS support before increasing this timeout value as it will impact Mover behavior in other ways, such as error and recovery scenarios.

/etc/services, /etc/inetd.conf, and /etc/xinetd.d

To invoke the remote part of the Mover, **inetd** on the remote node is utilized to start the parent process when a connection is made to a port based on the Mover’s type specific configuration (see *Section 5.2.6, “Mover-specific configuration”*).

An entry must be added to the `/etc/services` file so that **inetd** will listen on the port to which the Mover client process will connect. The port number is one greater than that configured for the "TCP Listen Port" in the Mover’s type-specific configuration. For example, if the configured listen port is 5001, the following line must be added to `/etc/services`:

```
hpss_mvrl 5002/tcp
```

For non-Linux systems. An entry must be added to the `/etc/inetd.conf` file which defines which executable to run and the arguments to use when a connection is detected. The entry must use the service name specified for the Mover in `/etc/services`. The sole argument (other than the program name) is the pathname to a file that contains an ASCII representation of the configured encryption key for the Mover (see the next section for further details). For example, for the `/etc/services` entry added above, the corresponding `/etc/inetd.conf` entry would be:

```
hpss_mvrl stream tcp nowait root /opt/hpss/bin/hpss_mvr_tcp  
hpss_mvr_tcp /var/hpss/etc/mvr_ek
```

This will cause **inetd** to run the executable `/opt/hpss/bin/hpss_mvr_tcp` under the root user ID when a connection is detected on port 5002. The Mover process uses the `/var/hpss/etc/mvr_ek` file to read the encryption key that will be used to authenticate all connections made to this Mover.

After modifying the `/etc/inetd.conf` file, be sure to refresh the `inetd` daemon using the following commands:

```
% ps -ef | grep inetd
root 6450 3154 0 Apr 29 - 0:02 /usr/sbin/inetd
hpss 17852 59370 2 16:50:25 pts/18 0:00 grep inetd
% kill -1 6450
```

For Linux systems. A file must be added to the `/etc/xinetd.d` directory which defines which program to run and the arguments to use when a connection is detected. The file will be given the same name as the service name specified for this Mover in `/etc/services`. For example, for the `/etc/services` entry added above, the corresponding file `/etc/xinetd.d/hpss_mvr1` would be created with the following contents:

```
service hpss_mvr1
{
    disable = no
    socket_type = stream
    protocol = tcp
    flags = IPv4
    wait = no
    port = 5002
    user = root
    server = /opt/hpss/bin/hpss_mvr_tcp
    server_args = -s /var/hpss/etc/mvr_ek
}
```

The specified port will be one greater than the port listed as the TCP Listen Port in the Mover's type-specific configuration. For example, the port value in the example corresponds to a Mover with a TCP Listen Port value of 5001.

This will cause **inetd** to run the executable `/opt/hpss/bin/hpss_mvr_tcp` under the root user ID when a connection is detected on port 5002. The Mover process will use the `/var/hpss/etc/mvr_ek` file to read the encryption key that will be used to authenticate all connections made to this Mover.

The **-s** option will enable syslog messages for startup debugging. This logging mechanism will be used prior to the initialization of the HPSS logging system and can be very useful in determining startup problems. The messages will be logged to the system log using the *user* facility and importance level of *alert*.

After modifying the file in `/etc/xinetd.d`, be sure to refresh `xinetd` daemon using the following commands:

```
% /sbin/service xinetd --full-restart
Stopping xinetd: [ OK ]
Starting xinetd: [ OK ]
```

The Mover encryption key files

To authenticate access to the remote Mover processes, the encryption key configured in this Mover's specific configuration (see *Section 5.2.6, "Mover-specific configuration"*) is read from a file accessible from the local file system. This file contains an ASCII representation of the encryption key. The pathname of the file is passed to the Mover executable as specified in either the `/etc/inetd.conf` or `/etc/xinetd.d` file. For example, if the encryption key in the Mover's type-specific

configuration is 1234567890ABCDEF, then the encryption key file (`/var/hpss/etc/ek.mvr1`) should contain:

```
0x12345678 0x90ABCDEF
```

/var/hpss/etc files required for remote Mover

The Mover process on the remote machine requires access to the following files in `/var/hpss/etc`:

- `authz.conf`
- `env.conf`
- `HPSS.conf`

If the remote Mover process will have access to SAN3p devices, the following file is required.

- `hpss_san3p.conf`

The remaining files should be copied from the HPSS root subsystem machine.

In addition to the above files the Mover process on the remote machine requires that the `/var/hpss/tmp` directory be present with write and search permissions for the root user.

See *The /var/hpss files* appendix of the *HPSS Installation Guide* for a description of each of these files.

System configuration parameters on Linux

This section describes a number of system configuration parameters which may need to be modified for the successful operation of the remote Mover.

Refer to the *Semaphore values* section of the *HPSS Installation Guide* (in the *System preparation* chapter) for detailed information on setting system parameters.

Setting up remote Movers with mkhpss

The **mkhpss** utility may be used to copy the files needed for a remote Mover from the root subsystem machine, to create the files which may not be copied, to install the required files on the remote Mover machine, and to configure the **inetd** to run the remote Mover.

To create the necessary files, run **mkhpss** on the root subsystem machine. From the **Root Subsystem Machine** menu, select the **Create Config Bundle** option. From this window, set the desired name of the config file bundle and click the **Create Config Bundle** button.

To install the files and configure the **inetd** on the remote Mover machine, transfer the created config bundle to the remote Mover machine and run **mkhpss** on the remote Mover machine. From the **Mover/Client Machine** menu, select the **Install Config Bundle** item. On this window, set the appropriate name of the bundle file and click the **Install Config Bundle** button.

Next, select the **Mover/Client Machine** menu **Configuration** submenu, and from this, the **Security Services** item. From this window, select the authentication mechanism (Kerberos or UNIX) using

the drop down menu of the **Authentication Mechanism** field. For Kerberos, complete the remaining fields on the window and click the **Configure Security Services** button to set up the required files. For UNIX, there is nothing further to configure on this window.

Next, still under the **Mover/Client Machine** menu **Configuration** submenu, select the **Other Services** item. On this window, select the check box for **Configure any Movers for this machine**. This check box instructs **mkhpss** to add the proper lines for the Movers to `/etc/services`. It also instructs it, depending upon the operating system type, either to modify the `/etc/inetd.conf` file or to add the appropriate configuration file to the `/etc/xinetd.d` directory to support the Movers. On this same window, select the check box for **Set any system parameters necessary for HPSS**; this configures the local system settings described in the section called “*System configuration parameters on Linux*”. At this time, there are no additional system parameters to set for AIX, so it is not actually necessary to select this check box for AIX hosts.

Mover configuration to support local file transfer

The Mover local file transfer (LFT) protocol allows the Movers to transfer data directly between a UNIX file system and HPSS.

Caveats:

- This transfer can only occur if both the Mover and the clients have direct access to the UNIX file system on which the file resides.
- If the multinode features of the Parallel FTP Client are used, the UNIX file system must be global to all nodes or unexpected results may occur.
- The Mover must be built with the LFT option. This is the default option for all Movers. If not all Movers have been built with this option, clients must explicitly specify a Class of Service which is valid for a Mover supporting the local file transfer option.
- The Mover must be running as root. If there are other Movers running on the same node, they must also run as root to take advantage of the Mover-to-Mover shared memory data transfer.
- A new configuration file (`/var/hpss/etc/hpss_mvr_localfilepath.conf`) is required on all nodes running an enhanced Mover.

The configuration file contains a list of UNIX paths. These paths specify which files are allowed to be moved using this special data transfer protocol. Any UNIX file on the Mover node whose path starts with a path in the list of entries is allowed to be transferred using the special data protocol.

If the configuration file does not exist, then the SSM will issue a minor alarm when the Mover starts, and no UNIX files will be transferred with the special protocol.

Here is an example configuration file:

```
# This file contains Unix paths that can be accessed by the
# local Movers on this node. If the client wishes to make a
# local file transfer and the start of the client path matches
# any of the paths in this list then the transfer will proceed,
# otherwise the Mover will not transfer the file.
```

```
#  
# The format of this file is simply a list of paths, one per  
# line.  
/gpfs  
/local/globalfilesystem
```

In the above sample configuration, any file under the path `/gpfs` or the path `/local/globalfilesystem` can be transferred using the special data protocol subject to the caveats specified above.

Extreme care should be taken when using this feature. The Mover will move data from any file in its list. If the UNIX file system is not global to all Mover nodes, the request will fail.



The following commands have been added to the HPSS Parallel FTP Clients to take advantage of these new features: `lfget`, `lfput`, `mlfget`, `mlfput`, `lfappend`. See the HPSS User's Guide for more details. The `hpss_ReadList` and `hpss_WriteList` Client APIs can also be used to utilize the local file transfer capability. See the HPSS Programmer's Reference - I/O Supplement for more details.

5.2.7. Physical Volume Repository (PVR)-specific configuration

The PVR-specific configuration entry can be created and managed using the *PVR Server Configuration* window. If changes are made while the server is running, the server must be recycled in order for the changes to become effective.

If you are configuring a PVR for StorageTek, ADIC AML, or generic SCSI: before proceeding with PVR configuration you should read the PVR-specific sections (Section 5.2.7.4, “STK PVR-specific configuration window”, Section 5.2.7.2, “AML PVR-specific configuration window”, and Section 5.2.7.3, “SCSI PVR-specific configuration window”). These sections provide additional vendor-specific advice on PVR/robot configuration.

The AML PVR is supported by special bid only.

5.2.7.1. Operator PVR-specific configuration window

Operator PVR Configuration

File Edit Help

Basic Controls

Server Name: Operator PVR

Server ID: 12

UUID: 8bcc78a6-8e1d-11e5-83b3-005056b7004f

Server Type: PVR

Server Subtype: Operator

Log Policy: Default

Execution Controls **Interface Controls** **Security Controls** **Audit Policy** **Specific**

Cartridge Capacity: 1000

Cartridge Alarm Threshold: 80 percent

Shelf Tape Check-In Retry: 30 seconds

Shelf Tape Check-In Alarm: 10 minutes

Dismount Delay: 15 minutes

Defer Dismount Exempt Count: -1

Characteristics

☒ Defer Dismounts ☐ Support Shelf Tape

Update Delete Start Over Clone (partial) Dismiss

Field descriptions.

Cartridge Capacity

The total number of cartridge slots in the library dedicated to this HPSS PVR. This may or may not be the total cartridge capacity of the library; a site might use part of the library for some other HPSS PVR or for some non-HPSS application. The PVR uses the **Cartridge Capacity** field and the **Cartridge Alarm Threshold** field to determine when to send an alarm that the total cartridge threshold has been exceeded.

Cartridge Alarm Threshold

The percentage of the **Cartridge Capacity** at which the PVR will send an alarm.

Shelf Tape Check-In Retry

The number of seconds the PVR will wait before asking the robot if a requested shelf tape has been checked in. The PVR will continue checking at this interval until the tape is checked in. This

field applies only if the **Support Shelf Tape** check box is selected. The retry value must be 30 or greater.

Shelf Tape Check-In Alarm

The PVR will periodically log alarm messages when a requested shelf tape has not been checked in. This field specifies the number of minutes between alarms. This field applies only if the **Support Shelf Tape** check box is selected. The alarm value must be "2" or greater.

Dismount Delay

When **Defer Dismounts** is checked, this value is used by the PVL to determine the number of minutes that dismounts are delayed after the last data access.

Defer Dismount Exempt Count

If set greater than "0", this value will specify the number of drives within a PVR that will be exempt from the defer dismount logic. This behavior will take place for all drives within the PVR with a drive/device "Library Unit ID" set (see Drive/Device Configuration). If multiple Library Unit IDs are set, the same Exempt count will apply to all Library Unit ID subsets.

The goal of this option is to allow a finer control of the defer dismount logic to account for transfer costs (for example, elevator, pass through costs). Defer dismount logic to apply logic on a sub library case (unit). The definition of a "unit" is placed in the hands of the administrator and likely based on library characteristics and topology.

Characteristics

Flags for the PVR:

Defer Dismounts

If ON, the PVL will delay the dismounting of a tape cartridge until the drive is required by another job or until the **Dismount Delay** time limit is exceeded. If the **Defer Dismount Exempt Count** is also set, the defer dismount logic will not be exercised on the number of drives set. This may result in free drives available for new mount requests without first generating defer dismount operations. See **Defer Dismount Exempt Count** for details.

A PVL restart is required in order for changes to this flag to take effect.

Support Shelf Tape

If ON, the PVR and the PVL will support the removal of cartridges from the tape library using the **shelf_tape** utility.

5.2.7.2. AML PVR-specific configuration window

The screenshot shows the 'AML PVR Configuration' window. The 'Basic Controls' tab is active, displaying fields for Server Name (AML PVR), Server ID (174), UUID (43c5f80e-311b-11e8-ad3c-005056b7004f), Server Type (PVR), Server Subtype (AML), and Log Policy (Default). Below this, the 'Specific' tab is selected, showing a list of configuration parameters with input fields and units. At the bottom, there are checkboxes for 'Defer Dismounts' and 'Support Shelf Tape', followed by fields for Client Name and Server Name. A row of buttons (Update, Delete, Start Over, Clone (partial), Dismiss) and a status message 'Successfully saved the 'AML PVR Configuration'' are at the very bottom.

Parameter	Value	Unit
Cartridge Capacity	3000	
Cartridge Alarm Threshold	90	percent
Same Job on Controller	10	
Other Job on Controller	5	
Distance To Drive	2	
Shelf Tape Check-In Retry	30	seconds
Shelf Tape Check-In Alarm	10	minutes
Dismount Delay	15	minutes
Retry Mount Time Limit	-1	minutes
Drive Error Limit	-1	
Defer Dismount Exempt Count	-1	

Characteristics

☐ Defer Dismounts ☐ Support Shelf Tape

Client Name

Server Name

Update Delete Start Over Clone (partial) Dismiss

Successfully saved the 'AML PVR Configuration'

Field descriptions.

Cartridge Capacity

The total number of cartridge slots in the library dedicated to this HPSS PVR. This may or may not be the total cartridge capacity of the library; a site might use part of the library for some other HPSS PVR or for some non-HPSS application. The PVR uses the **Cartridge Capacity** field and

the **Cartridge Alarm Threshold** field to determine when to send an alarm that the total cartridge threshold has been exceeded.

Cartridge Alarm Threshold

The percentage of the **Cartridge Capacity** at which the PVR will send an alarm.

Same Job on Controller, Other Job on Controller, Distance To Drive

These three values are used by the PVR when selecting a drive for a tape mount operation. The three values are essentially weights that are used to compute an overall score for each possible drive. After the score has been calculated, the drive with the lowest score is selected for the mount. If two or more drives tie for the lowest score, one drive is selected at random. The score is calculated as follows:

$$\begin{aligned} \text{Score} = & \\ & \text{Weight 1} \times \text{Cartridges from this job mounted on this drive's controller} + \\ & \text{Weight 2} \times \text{Cartridges from other jobs mounted on this drive's controller} + \\ & \text{Weight 3} \times \text{Units of distance from the cartridge to the drive} \end{aligned}$$

This method has the effect of distributing a striped tape mount across as many controllers as possible for the best performance. It also will try to pick controllers that are currently driving a minimum number of tapes. So, in an environment with many tape drives per controller, the best performance will be achieved by minimizing the load on any one controller.

The **Distance To Drive** helps minimize mount times by mounting the tape in a physically close drive. All other things being equal, the tape will be mounted in the closest drive. However, if the closest drive is attached to a heavily used controller, then a more distant drive will be selected.

Shelf Tape Check-In Retry

The number of seconds the PVR will wait before asking the robot if a requested shelf tape has been checked in. The PVR will continue checking at this interval until the tape is checked in. This field applies only if the **Support Shelf Tape** check box is selected. The retry value must be "30" or greater.

Shelf Tape Check-In Alarm

The PVR will periodically log alarm messages when a requested shelf tape has not been checked in. This field specifies the number of minutes between alarms. This field applies only if the **Support Shelf Tape** check box is selected. The alarm value must be "2" or greater.

Dismount Delay

When **Defer Dismounts** is checked, this value is used by the PVL to determine the number of minutes that dismounts are delayed after the last data access.

Retry Mount Time Limit

The default value for this field is "-1". When the default value is used, if an error is encountered during a PVR mount operation, the mount will pend and be retried every five minutes. Setting a value in this field will change the mount behavior to periodically retry the mount until the specified time limit is exceeded. Once exceeded, an error is generated and the mount request is canceled. If the PVR mount operation is for a RAIT volume, some number of mount failures can occur and still allow the mount to complete successfully. When the PVR is mounting RAIT volumes, it is not recommended to set this field to "-1" because this setting will make a RAIT PVL

mount job wait up to 24 hours for all tapes to be mounted. It is recommended to set this field to a small time limit, such as 30 minutes, so that RAIT PVL mount jobs will not be blocked for such a long time. If the mount request would have resulted in a write operation, the error returned will cause the Core Server to set the **VV Condition** of the associated tape volume to DOWN. Once in DOWN state, the volume will no longer be available for read or write operations. For further information about the Core Server VV Condition, see *Section 8.5.4.2, "Core Server Tape Volume information window"*.

Drive Error Limit

This field is intended to be used in conjunction with the PVR Server **Retry Mount Time Limit**. If the number of consecutive mount errors which occur to any drive in this PVR equal or exceed this value, the drive is automatically locked by the PVL. The only mount errors that apply are those set through the **Retry Mount Time Limit** mechanism. The **Drive Error Count** field in the *PVL Drive Information* records the number of consecutive errors on a drive by drive basis. To turn off the automatic drive disable feature, set the **Drive Error Limit** to "0" or "-1". Since this field is employed by the PVL, changing its value will not change drive disable behavior until the PVL is recycled.

Defer Dismount Exempt Count

If set greater than "0", this value will specify the number of drives within a PVR that will be exempt from the defer dismount logic. This behavior will take place for all drives within the PVR with a drive/device "Library Unit ID" set (see Drive/Device Configuration). If multiple Library Unit IDs are set, the same Exempt count will apply to all Library Unit ID subsets.

The goal of this option is to allow a finer control of the defer dismount logic to account for transfer costs (for example, elevator, pass through costs). Defer dismount logic to apply logic on a sub library case(unit). The definition of a "unit" is placed in the hands of the administrator and likely based on library characteristics and topology.

Characteristics

Flags for the PVR:

Defer Dismounts

If ON, the PVL will delay the dismounting of a tape cartridge until the drive is required by another job or until the **Dismount Delay** time limit is exceeded. If the **Defer Dismount Exempt Count** is also set, the defer dismount logic will not be exercised on the number of drives set. This may result in free drives available for new mount requests without first generating defer dismount operations. See **Defer Dismount Exempt Count** for details.

A PVL restart is required in order for changes to this flag to take effect.

Support Shelf Tape

If ON, the PVR and the PVL will support the removal of cartridges from the tape library using the **shelf_tape** utility.

Client Name

The name of the client requesting authorization from the Distributed Automated Media Library Server. DAS software, which executes on the OS/2 controller PC, allows different clients to control the AML robotics system. DAS uses the client name to determine the access permission of the requesting client to the AML's storage positions, drives, and Insert/Eject units. Access

configurations for clients are set in the configuration file `C:\DAS\ETC\CONFIG` on the OS/2 PC. The client name can be up to 64 alphanumeric characters in length and is case sensitive.

Server Name

TCP/IP hostname or IP address of the AML OS/2-PC DAS server. This value must be defined in the network domain server and must be resolvable during DAS start. The server name is set in the configuration file `C:\CONFIG.SYS` on the OS/2 PC. The server name can be up to 64 alphanumeric characters long and can include up to six dots (.).

AML PVR additional information

The AML PVR is supported by special bid only.

ADIC Automatic Media Library storage systems information.

HPSS can share an AML robot with other tape management systems. If a robot is shared, care must be taken to make sure that a drive is not used by any other tape management system while that drive is configured as unlocked in HPSS. This is important because HPSS can be configured to periodically poll all of its unlocked drives even if they are not currently mounted or in use by HPSS. If a drive is being used by another tape management system, it must be configured as locked in HPSS. For robots that have more than one arm (such as the AML/2), users should configure the PVL Drive Information/Controller ID of each drive depending on which arm is servicing it.

User needs to set the Server Name and Client Name, which are case sensitive, in the AML PVR Server Configuration panel to establish the connectivity between the HPSS software and the OS/2 controlling the robot. The Server Name is the name of the controller associated with the TCP/IP address, as defined in the TCP/IP HOST file, and the Client Name is the name of the OS/2 administrator client as defined in the DAS configuration.

Additionally users must configure the Insert/Eject ports for the AML PVR using the configuration files `/var/hpss/etc/AML_EjectPort.conf` and `/var/hpss/etc/AML_InsertPort.conf`. The reasons are that the AML robot can have multiple insert and eject ports, which have the capability to handle different media types. These two configuration files in conjunction with the AMU AMS configuration files specify which Insert/Eject areas or bins the tapes should be placed in for insertion into the archive and where they are ejected to when the HPSS export command is used.

Note that if multiple E/I/F bins are used for exporting cartridges, users must set the environment variable `DAS_EJECTAREAFULL=1` on the DAS Server PC by doing the following:

1. Edit the `config.sys`
2. Add a new line:

```
set DAS_EJECTAREAFULL=1
```

Starting the DAS server processes. In order for the HPSS AML PVR to communicate with the AML robot, it is required that the DAS Server process be running.

Once the DAS, TCP/IP, and AML configuration have been completed, the DAS server is initialized and started with the following steps:

1. Make sure the AMU archive management software is running and the hostname is resolved,
2. Select an OS/2 window from the Desktop and change the directory to C:\DAS,

```
C:> cd \das
```

3. At the prompt, type "tcpstart" and make sure that TCP/IP gets configured and that the port mapper program is started,

```
C:\das> tcpstart
```

4. Type "dasstart" to start the DAS server

```
C:\das> dasstart
```

In most cases, a startup file like the following can be used to automatically start the DAS processor:

```
call tcpstart
das\tools\os2sleep 20
CD \AMU
START CON
START KRN
cd \das
tools\os2sleep 20
call dasstart
cd bin
start "DAS/2 AmuClient"
exit
```

5.2.7.3. SCSI PVR-specific configuration window

SCSI PVR Configuration

File Edit Help

Basic Controls

Server Name: SCSI PVR

Server ID: 17

UUID: 939835a4-9ebf-11e5-95ea-005056b7004f

Server Type: PVR

Server Subtype: SCSI

Log Policy: Alarms/Events Only

Execution Controls **Interface Controls** **Security Controls** **Audit Policy** **Specific**

Cartridge Capacity: 440

Cartridge Alarm Threshold: 90 percent

Same Job on Controller: 10

Other Job on Controller: 5

Distance To Drive: 0

Shelf Tape Check-In Retry: 30 seconds

Shelf Tape Check-In Alarm: 10 minutes

Dismount Delay: 15 minutes

Retry Mount Time Limit: -1 minutes

Drive Error Limit: -1

Defer Dismount Exempt Count: -1

Characteristics

☒ Defer Dismounts ☐ Support Shelf Tape ☐ Enforce Home Location

Serial Number 0: T10: IBM| 3584 78A3467040

Serial Number 1:

Serial Number 2:

Serial Number 3:

Serial Number 4:

Serial Number 5:

Serial Number 6:

Serial Number 7:

Update Delete Start Over Clone (partial) Dismiss

Field descriptions.**Cartridge Capacity**

The total number of cartridge slots in the library dedicated to this HPSS PVR. This may or may not be the total cartridge capacity of the library; a site might use part of the library for some other HPSS PVR or for some non-HPSS application. The PVR uses the **Cartridge Capacity** field and the **Cartridge Alarm Threshold** field to determine when to send an alarm that the total cartridge threshold has been exceeded.

Cartridge Alarm Threshold

The percentage of the **Cartridge Capacity** at which the PVR will send an alarm.

Same Job on Controller, Other Job on Controller, Distance To Drive

These three values are used by the PVR when selecting a drive for a tape mount operation. The three values are essentially weights that are used to compute an overall score for each possible drive. After the score has been calculated, the drive with the lowest score is selected for the mount. If two or more drives tie for the lowest score, one drive is selected at random. The score is calculated as follows:

$$\begin{aligned} \text{Score} = & \\ & \text{Weight 1} \times \text{Cartridges from this job mounted on this drive's controller} + \\ & \text{Weight 2} \times \text{Cartridges from other jobs mounted on this drive's controller} + \\ & \text{Weight 3} \times \text{Units of distance from the cartridge to the drive} \end{aligned}$$

This method has the effect of distributing a striped tape mount across as many controllers as possible for the best performance. It also will try to pick controllers that are currently driving a minimum number of tapes. So, in an environment with many tape drives per controller, the best performance will be achieved by minimizing the load on any one controller.

The **Distance To Drive** helps minimize mount times by mounting the tape in a physically close drive. All other things being equal, the tape will be mounted in the closest drive. However, if the closest drive is attached to a heavily used controller, then a more distant drive will be selected.

Shelf Tape Check-In Retry

The number of seconds the PVR will wait before asking the robot if a requested shelf tape has been checked in. The PVR will continue checking at this interval until the tape is checked in. This field applies only if the **Support Shelf Tape** check box is selected. The retry value must be "30" or greater.

Shelf Tape Check-In Alarm

The PVR will periodically log alarm messages when a requested shelf tape has not been checked in. This field specifies the number of minutes between alarms. This field applies only if the **Support Shelf Tape** check box is selected. The alarm value must be "2" or greater.

Dismount Delay

When **Defer Dismounts** is checked, this value is used by the PVL to determine the number of minutes that dismounts are delayed after the last data access.

Retry Mount Time Limit

The default value for this field is "-1". When the default value is used, if an error is encountered during a PVR mount operation, the mount will pend and be retried every five minutes. Setting a value in this field will change the mount behavior to periodically retry the mount until the

specified time limit is exceeded. Once exceeded, an error is generated and the mount request is canceled. If the mount request would have resulted in a write operation, the error returned will cause the Core Server to set the **VV Condition** of the associated tape volume to DOWN. Once in DOWN state, the volume will no longer be available for read or write operations. For further information about the Core Server VV Condition, see *Section 8.5.4.2, "Core Server Tape Volume information window"*.

Drive Error Limit

This field is intended to be used in conjunction with the PVR Server **Retry Mount Time Limit**. If the number of consecutive mount errors which occur to any drive in this PVR equal or exceed this value, the drive is automatically locked by the PVL. The only mount errors that apply are those set through the **Retry Mount Time Limit** mechanism. The **Drive Error Count** field in the *PVL Drive Information* records the number of consecutive errors on a drive by drive basis. To turn off the automatic drive disable feature, set the **Drive Error Limit** to "0" or "-1". Since this field is employed by the PVL, changing its value will not change drive disable behavior until the PVL is recycled.

Defer Dismount Exempt Count

If set greater than "0", this value will specify the number of drives within a PVR that will be exempt from the defer dismount logic. This behavior will take place for all drives within the PVR with a drive/device "Library Unit ID" set (see Drive/Device Configuration). If multiple Library Unit IDs are set, the same Exempt count will apply to all Library Unit ID subsets.

The goal of this option is to allow a finer control of the defer dismount logic to account for transfer costs (for example, elevator, pass through costs). Defer dismount logic to apply logic on a sub library case (unit). The definition of a "unit" is placed in the hands of the administrator and likely based on library characteristics and topology.

Characteristics

Flags for the PVR:

Defer Dismounts

If ON, the PVL will delay the dismounting of a tape cartridge until the drive is required by another job or until the **Dismount Delay** time limit is exceeded. If the **Defer Dismount Exempt Count** is also set, the defer dismount logic will not be exercised on the number of drives set. This may result in free drives available for new mount requests without first generating defer dismount operations. See **Defer Dismount Exempt Count** for details.

A PVL restart is required in order for changes to this flag to take effect.

Support Shelf Tape

If ON, the PVR and the PVL will support the removal of cartridges from the tape library using the **shelf_tape** utility.

Enforce Home Location

If ON, the SCSI PVR will always try to dismount a mounted cart back to its home location. Otherwise, it will just use the first free slot. The **scsi_home** utility can be used to view the home location values.

Serial Number 0 - 7

The serial number of the robots, obtained from **device_scan**. Only Serial Number 0 should be used. The serial number will allow the SCSI PVR to automatically look up all available control

paths upon startup. Below is an example output from the **device_scan** tool which shows the **library serial number** as well as the **drive serial number**. The **drive serial number** is used when configuring the library's drives and should be placed in the "Drive Address" field.

```
1) Device=[/dev/sg5]
   Vendor = [IBM]
   Product = [03584L22]
   Revision = [B071]
   Vendor info = [B071780000078A3467 0]
   Vital Product =
       Identifier = ['T10:IBM|03584L22          0000078A34670401'] // (Library serial number)
       Data Transfer Elements
           Address Serial
       1 = [257] ['T10:IBM|03592E06          0000078A426A'] // (Drive serial number)
       2 = [258] ['T10:IBM|03592E06          0000078A4276']
       3 = [259] ['T10:IBM|03592E06          0000078A4294']
       4 = [260] ['T10:IBM|03592E06          0000078A424D']
       5 = [261] ['T10:IBM|03592E06          0000078A4292']
       6 = [262] ['T10:IBM|03592E06          0000078A427A']
       7 = [263] ['T10:IBM|03592E06          0000078A4286']
       8 = [264] ['T10:IBM|03592E06          0000078A4261']
```

Note that everything between the single quotes should be included, including the internal spaces. The simplest way to do this is by copying and pasting the output into SSM.

5.2.7.4. STK PVR-specific configuration window

The screenshot shows the 'STK PVR Configuration' window. The 'Basic Controls' tab is active, displaying fields for Server Name (STK PVR), Server ID (175), UUID (284c0bf8-311c-11e8-ad3c-005056b7004f), Server Type (PVR), Server Subtype (STK), and Log Policy (Default). Below this is a tabbed interface with 'Execution Controls', 'Interface Controls', 'Security Controls', 'Audit Policy', and 'Specific'. The 'Specific' tab is selected, showing a list of configuration parameters: Cartridge Capacity (6000), Cartridge Alarm Threshold (90 percent), Same Job on Controller (10), Other Job on Controller (5), Distance To Drive (20), Shelf Tape Check-In Retry (30 seconds), Shelf Tape Check-In Alarm (10 minutes), Dismount Delay (15 minutes), Retry Mount Time Limit (-1 minutes), Drive Error Limit (-1), and Defer Dismount Exempt Count (-1). Below these is a 'Characteristics' section with checkboxes for 'Defer Dismounts' and 'Support Shelf Tape'. At the bottom, there are fields for 'ACSLs Packet Version' (4) and 'Alternate SSI Port'. A status bar at the very bottom indicates 'Successfully saved the 'STK PVR Configuration''.

Field	Value
Server Name	STK PVR
Server ID	175
UUID	284c0bf8-311c-11e8-ad3c-005056b7004f
Server Type	PVR
Server Subtype	STK
Log Policy	Default
Cartridge Capacity	6000
Cartridge Alarm Threshold	90 percent
Same Job on Controller	10
Other Job on Controller	5
Distance To Drive	20
Shelf Tape Check-In Retry	30 seconds
Shelf Tape Check-In Alarm	10 minutes
Dismount Delay	15 minutes
Retry Mount Time Limit	-1 minutes
Drive Error Limit	-1
Defer Dismount Exempt Count	-1
ACSLs Packet Version	4
Alternate SSI Port	

Field descriptions.

Cartridge Capacity

The total number of cartridge slots in the library dedicated to this HPSS PVR. This may or may not be the total cartridge capacity of the library; a site might use part of the library for some other HPSS PVR or for some non-HPSS application. The PVR uses the **Cartridge Capacity** field and

the **Cartridge Alarm Threshold** field to determine when to send an alarm that the total cartridge threshold has been exceeded.

Cartridge Alarm Threshold

The percentage of the **Cartridge Capacity** at which the PVR will send an alarm.

Same Job on Controller, Other Job on Controller, Distance To Drive

These three values are used by the PVR when selecting a drive for a tape mount operation. The three values are essentially weights that are used to compute an overall score for each possible drive. After the score has been calculated, the drive with the lowest score is selected for the mount. If two or more drives tie for the lowest score, one drive is selected at random. The score is calculated as follows:

$$\begin{aligned} \text{Score} = & \\ & \text{Weight 1} \times \text{Cartridges from this job mounted on this drive's controller} + \\ & \text{Weight 2} \times \text{Cartridges from other jobs mounted on this drive's controller} + \\ & \text{Weight 3} \times \text{Units of distance from the cartridge to the drive} \end{aligned}$$

This method has the effect of distributing a striped tape mount across as many controllers as possible for the best performance. It also will try to pick controllers that are currently driving a minimum number of tapes. So, in an environment with many tape drives per controller, the best performance will be achieved by minimizing the load on any one controller.

The **Distance To Drive** helps minimize mount times by mounting the tape in a physically close drive. For STK drives, a unit of distance is from one Silo/LSM (Library Storage Module) to the next. This means that the tape must go through a pass-through port or elevator. This process is more time consuming. Therefore the default value for STK robots is a higher value which forces the tape to stay in the same Silo/LSM even if it means the drive selected is attached to a controller which is heavily used. All other things being equal, the tape will be mounted in the closest drive.

Shelf Tape Check-In Retry

The number of seconds the PVR will wait before asking the robot if a requested shelf tape has been checked in. The PVR will continue checking at this interval until the tape is checked in. This field applies only if the **Support Shelf Tape** check box is selected. The retry value must be "30" or greater.

Shelf Tape Check-In Alarm

The PVR will periodically log alarm messages when a requested shelf tape has not been checked in. This field specifies the number of minutes between alarms. This field applies only if the **Support Shelf Tape** check box is selected. The alarm value must be "2" or greater.

Dismount Delay

When **Defer Dismounts** is checked, this value is used by the PVL to determine the number of minutes that dismounts are delayed after the last data access.

Retry Mount Time Limit

The default value for this field is "-1". When the default value is used, if an error is encountered during a PVR mount operation, the mount will pend and be retried every five minutes. Setting a value in this field will change the mount behavior to periodically retry the mount until the specified time limit is exceeded. Once exceeded, an error is generated and the mount request is canceled. If the mount request would have resulted in a write operation, the error returned will cause the Core Server to set the **VV Condition** of the associated tape volume to DOWN. Once

in DOWN state, the volume will no longer be available for read or write operations. For further information about the Core Server VV Condition, see *Section 8.5.4.2, “Core Server Tape Volume information window”*.

Drive Error Limit

This field is intended to be used in conjunction with the PVR Server **Retry Mount Time Limit**. If the number of consecutive mount errors which occur to any drive in this PVR equal or exceed this value, the drive is automatically locked by the PVL. The only mount errors that apply are those set through the **Retry Mount Time Limit** mechanism. The **Drive Error Count** field in the *PVL Drive Information* records the number of consecutive errors on a drive by drive basis. To turn off the automatic drive disable feature, set the **Drive Error Limit** to "0" or "-1". Since this field is employed by the PVL, changing its value will not change drive disable behavior until the PVL is recycled.

Defer Dismount Exempt Count

If set greater than "0", this value will specify the number of drives within a PVR that will be exempt from the defer dismount logic. This behavior will take place for all drives within the PVR with a drive/device "Library Unit ID" set (see Drive/Device Configuration). If multiple Library Unit IDs are set, the same Exempt count will apply to all Library Unit ID subsets.

The goal of this option is to allow a finer control of the defer dismount logic to account for transfer costs (for example, elevator, pass through costs). Defer dismount logic to apply logic on a sub library case (unit). The definition of a "unit" is placed in the hands of the administrator and likely based on library characteristics and topology.

Characteristics

Flags for the PVR:

Defer Dismounts

If ON, the PVL will delay the dismounting of a tape cartridge until the drive is required by another job or until the **Dismount Delay** time limit is exceeded. If the **Defer Dismount Exempt Count** is also set, the defer dismount logic will not be exercised on the number of drives set. This may result in free drives available for new mount requests without first generating defer dismount operations. See **Defer Dismount Exempt Count** for details.

A PVL restart is required in order for changes to this flag to take effect.

Support Shelf Tape

If ON, the PVR and the PVL will support the removal of cartridges from the tape library using the **shelf_tape** utility.

ACSLS Packet Version

The packet version used by STK's ACSLS software.

Alternate SSI Port

A site running two SSI (Server System Interface) services on the same platform will specify the non-default port number of the second service via this field.

STK PVR additional information

When configuring HPSS to manage an STK Robot, the Drive Address configuration entries correspond to the ACS, Unit, Panel, and Drive number used by ACSLS to identify drives. For example, the first drive in a typical STK Robot configuration has the Drive Address (0,0,10,0).

HPSS can share an STK Robot with other tape management systems. If a robot is shared, care must be taken to make sure that a drive is not used by any other tape management system while that drive is configured as unlocked in HPSS. This is important because HPSS can be configured to periodically poll all of its unlocked drives even if they are not currently mounted or in use by HPSS. If a drive is being used by another tape management system, it must be configured as locked in HPSS.

HPSS will use any Cartridge Access Port (CAP) in the STK Robot that has a priority greater than zero. When it needs a CAP, HPSS will pick the highest priority CAP that is currently available. At least one CAP must be assigned a nonzero priority. See the *STK Automated Cartridge System Library Software (ACSL) System Administrator's Guide* for procedures to set CAP priority.

ACSL Packet Version.

See the *STK ACSL System Administrator's Guide* for details. The environment variable `ACSAPI_PACKET_VERSION` will override the value entered in the **ACSL Packet Version** field of the *STK PVR-specific configuration window*. If neither are set, a default value of "4" is used since ACSL 7.x and 8.x generally use packet version 4.

Starting the STK client processes.

In order for the HPSS PVR to communicate with the STK robot, it is required that STK client-side processes be running on the node where the HPSS PVR is executing. These client-side processes are the Server System Interface (SSI) and the Toolkit event logger. These binaries and associated script files are maintained by Oracle, but can be obtained from HPSS support.

The deliverable includes the binaries and script files necessary for starting the STK client-side processes as well as documentation for these files. Binaries are available for the AIX, Linux x86_64, and Linux PPC platforms.

The `t_startit.sh` file must be executed to start the STK client processes on those nodes where the HPSS PVR is executing. The script will prompt for the following options:

1. Multi-Host or Single-Host: specify "m" for multiple host.
2. Server side or Client side: specify "c" for client-side.
3. Remote Host Name: specify the name of the host where the ACSL software is running.
4. Real ACSL server or simulated test: specify "r" for real ACSL server.
5. Remote Host Version: specify the remote host version which is listed for the ACSL Release Number. This will usually be "4":

ACSL Release Number	Remote Host Version Number
5.x, 6.x, 7.x, 8.x	4

6. Whether processes created should be listed (Y or N): specify either "y" or "n".
7. Whether `t_driver` should be started: specify "n".

To terminate the STK client processes, the script `t_killit.sh` may be executed.

Sample output from `t_startit.sh` follows:

```

***** Welcome to t_startit.sh, Version 2.3 *****
This is the automated tester and startit script for the TOOLKIT.*
Simply answer the questions which follow.*
Executables live in: /opt/hpss/stk/bin*
Would you like Multi-Host or Single-Host testing?*
Enter one of the following followed by ENTER:*
M Multi-host testing*
S Single-host testing*
X eXit this script*
Enter choice: m*
Would you like to define the server side or client side*
for Multi-Host testing?*
Enter one of the following followed by ENTER:*
S Server side*
C Client side*
Enter choice: c*
The Remote Host Name is the name of the server which*
has the ACSLS software (or simulator) running on it.*
Enter Remote Host Name (CSI_HOSTNAME): jeep*
Is that host running a real ACSLS (or LibStation) server,*
or a Simulated Test server (t_acslm)?*
Enter one of the following followed by ENTER:*
R Real ACSLS or LibStation Server*
S Simulated Test Server*
Enter choice: r*
The Remote Host Version number is the ACSLS Packet*
Version level which the server is expecting.*
Here are the valid choices:*
ACSLS Release Number      Remote Host Version Number*

        3.x                2*
        4.x                3*
        5.x, 6.x, 7.x, 8.x  4*
Enter Remote Host Version (ACSAPI_PACKET_VERSION): 4*
Starting /opt/hpss/stk/bin/mini_el...*
Attempting startup of /opt/hpss/bin/mini_el ...*
Starting /opt/hpss/bin/ssi...*
Attempting startup of PARENT for /opt/hpss/bin/ssi...*
SIGHUP received Parent Process ID is: 17290*
Attempting startup of /opt/hpss/bin/ssi...*
SIGHUP received Parent Process #17290 EXITING NORMALLY*
Initialization Done.*
Do you want to see the processes created? (Y or N): y*
17288 p4 S 0:00 /opt/hpss/stk/bin/mini_el*
17292 p4 S 0:00 /opt/hpss/stk/bin/ssi 17290 50004 23*
17295 p4 S 0:00 grep /opt/hpss/stk/bin*
Do you want to start up t_cdriver? (Y or N): n*

```

IPv6 limitation

The HPSS STK PVR is not supported on IPv6 due to the Oracle StorageTek CDK ACSAPI library not supporting it at this time. If you require this support, consult your Oracle representative about implementing this enhancement.

5.2.8. RAIT Engine-specific configuration

5.2.8.1. RAIT Engine-specific configuration window

RAIT Engine Configuration

File Edit Help

Basic Controls

Server Name: RAIT Engine (dus02 Remote)

Server ID: 162

UUID: 6c1d7d9e-3a55-11e7-b63d-005056b7004f

Server Type: RAIT Engine

Log Policy: Default

Execution Controls **Interface Controls** **Security Controls** **Audit Policy** **Specific**

Encryption Key: 0x9ef9b7c1d5d9f1fc [Generate](#)

TCP Path Name: /opt/hpss/bin/hpss_rait_engine_tcp

Hostname: hpss-dev-dus02.ccs.ornl.gov

TCP Port: 5031

Port Range Start: 0

Port Range End: 0

Data Hostnames

0: hpss-dev-dus02.ccs.ornl.gov

1:

2:

3:

4:

5:

6:

7:

[Update](#) [Delete](#) [Start Over](#) [Clone \(partial\)](#) [Dismiss](#)

This tab of the *RAIT Engine Configuration* window allows you to update and view the type-specific configuration for the RAIT Engine.

The server must be recycled in order for changes to become effective.

Field descriptions.

Encryption Key

An encryption key used to secure the RAIT Engine's interface with the Core Server. A specific value may be entered in the field or a random value may be generated by clicking on the **Generate** button. The default value is "0". A nonzero value must be configured to allow client access to the RAIT engine's data interface.

TCP Path Name

The pathname of the RAIT Engine TCP/IP process. The default value is `"/opt/hpss/bin/hpss_rait_engine_tcp"`.

Hostname

The name of the host interface used for RAIT Engine control communication. This must be a valid hostname for one of the network interfaces on the RAIT Engine node. The **Execute Hostname** of the *Core Server Configuration* window is used as the default. If the RAIT Engine is running remotely, this field must correspond to a network interface on the remote node, whereas the **Execute Hostname** set in the RAIT Engine's basic configuration corresponds to a network interface on which the administrative portion of the RAIT Engine runs.

Data Hostnames

The host network interface names to be used by the RAIT Engine when transferring data. There can be up to eight data host names specified. These must be valid hostnames for network interfaces on the RAIT Engine node. The default value of the first hostname is the **Execute Hostname** of the *Core Server Configuration* window. This field must correspond to a network interface on the remote node. These network interface addresses are utilized during migration/stage to allow the RAIT Engine to distribute connections and data across all the available network interfaces.

TCP Port

The TCP/IP port number used by the administrative interface portion of the RAIT Engine for the TCP/IP listen process to receive connections. It should be a value over 6000. The default value is 6001. The port number must be unique for processes running on the same node.

The RAIT Engine will use a port one greater than the value in this field on the remote node during initialization (for example, if 6001 is entered, port 6002 is used on the remote node); therefore available port ranges on both nodes must be taken into consideration when selecting this value.

Port Range Start

The beginning of a range of local TCP/IP port numbers to be used by the RAIT Engine when connecting to clients. Valid values are zero or any valid TCP port number to which the RAIT Engine may bind, that is less than or equal to the value of **Port Range End**. The default value is "0". Use of particular port ranges may be required by some sites for communication across a firewall. If this value is set to zero, the operating system will select the port number; otherwise the RAIT Engine selects a local port number between **Port Range Start** and **Port Range End** inclusive. If nonzero, this field must be less than or equal to **Port Range End**. If this field is zero, **Port Range End** field must also be zero. Note that in order for port range configuration to have an effect, client processes have to have the client API variable `HPSS_API_USE_PORT_RANGE` set (`HPSS_API_USE_PORT_RANGE=1`).

Port Range End

Used in conjunction with **Port Range Start** (see above). Valid values are zero or any TCP port number to which the RAIT Engine may bind (that is greater than or equal to the value of **Port Range Start**). The default value is "0". If nonzero, this field must be equal or greater than **Port Range Start**. If this field is zero, **Port Range Start** field must also be zero.

5.2.9. Deleting a server configuration

A server's configuration can be removed when the server is no longer needed. However, care should be taken to ensure that all objects created by the server are deleted properly and that all references to the server are removed from the HPSS configuration.



The steps described in this section are general guidelines. Specific procedures should be worked out with the aid of HPSS support so that the details of the system's configuration can be considered.

A server's configuration should be removed only when it is no longer needed. To modify a server's configuration, update the existing configuration instead of deleting the configuration and creating a new one. This will reduce the possibility of introducing configuration errors into HPSS.

There must be at least one Core Server, PVL, and Mover configuration entry in order for HPSS to be operational.

It is very important that the server's configuration entries be deleted in the order shown below.

To delete a server, perform the following steps in the order shown:

1. Delete all objects created by the server.
 - If deleting a Core Server configuration, realize that this should only be done when deleting the full subsystem.
 - If deleting a PVL configuration, all imported cartridges must first be exported.
 - If deleting a PVR configuration, all injected cartridges must either be exported or moved to another PVR.
 - If deleting a Gatekeeper, remove all references to that Gatekeeper from every subsystem configuration.
2. Remove references to the server from other configurations.



Do not delete Root Core Server configuration.

- If deleting a PVR configuration, update the **PVR** field in the appropriate device and drive configuration using *Section 7.1.1, "Devices and Drives window"*.
- If deleting a Mover configuration, update the Mover field in the appropriate device and drive configuration using *Section 7.1.1, "Devices and Drives window"*.

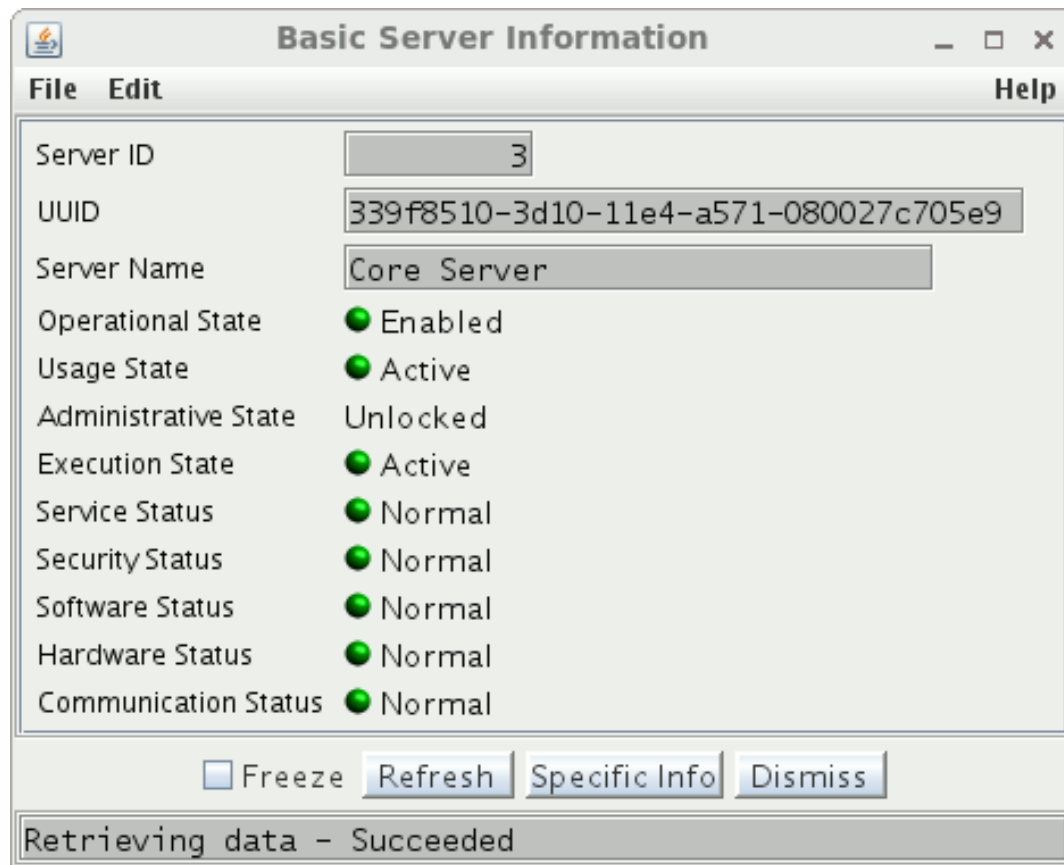
3. After the affected configurations are modified, reinitialize (if supported) or recycle the appropriate HPSS servers.
4. Ensure that the server is not running.
5. Delete UNIX files created or used by the server, if no longer needed.
 - If deleting the MPS, delete the MPS reports.
 - If deleting the Startup Daemon, delete all server lock files in the `/var/hpss/tmp` directory. The names of these files always begin with `hpssd`.
 - If deleting the Mover, delete Mover's socket files in `/var/hpss/tmp` directory.
6. Delete the server configuration by opening the server list, highlighting the target server, and clicking the **Delete** button.

5.3. Monitoring server information

A server that is running and connected to SSM will allow the SSM user to view and update its information. This section describes the server execution statuses and configuration information.

A typical HPSS server allows the SSM users to control its execution and monitor its server-related data through the *Basic Server Information* window and the server-specific information windows. These windows are described in the following subsections.

5.3.1. Basic Server Information



This window allows you to view the basic information associated with an HPSS server. The information is acquired from the server itself. While this window remains open, it will be automatically updated when the status of the server changes.

The *Basic Server Information* window is not available for viewing if the server's operational state is *Broken*, *Not Enabled*, or *Unknown*, or the System Manager is unable to contact the server.

Field descriptions.

Server ID

A unique value automatically generated and displayed for servers.

UUID

The universal unique identifier for the server.

Server Name

The descriptive name of the server.

Operational State

The operational state of the server. This state indicates the overall health of the server. Possible values are:

- *Enabled* indicates the server is running normally.
- *Disabled* indicates the server is not running because it has been shut down by SSM. You will not normally see this state on this window since the server must be running to provide the state information.
- *Suspect* indicates the server has seen an error that may indicate a problem is developing.
- *Minor* indicates the server has seen a problem which does not seriously impact HPSS operation.
- *Major* indicates the server has seen a problem which degrades HPSS operation.
- *Broken* indicates the server has shut itself down due to a serious error.

Usage State

The usage state of the server. Possible values are:

- *Active* indicates the server is working normally.
- *Idle* indicates the server is running, but has no work to perform. Most servers do not update usage state dynamically, so it is unlikely you will see this value reported.
- *Busy* indicates the server is busy performing its function. Most servers do not update usage state dynamically, so it is unlikely you will see this value reported.
- *Unknown* indicates the server has not reported a recognized usage state.

Administrative State

The Administrative State of the server. The possible states are:

- *Shut Down* indicates the server shut itself down in an orderly way.

- `Force Halt` indicates the server accepted a Force Halt command and terminated immediately.
- `Reinitialize` indicates the server reinitialized itself. Not all HPSS servers support being reinitialized.
- `Mark Repaired` indicates the server cleared its error states.
- `Unlocked` is displayed during normal server operation.

Execution State

The execution state of the server. Possible values are: `Active`, `Initializing`, `Restarting`, `Terminating`, `Inactive`, `Suspending`, `Suspended`, and `Resuming`.

Service Status

The server's service status. Possible values are: `Normal`, `Warning`, `Minor Problem`, `Major Problem`, and `Critical`.

Security Status

The server's security status. Possible values are: `Normal`, `Warning`, `Minor Problem`, `Major Problem`, and `Critical`.

Software Status

The server's software status. Possible values are: `Normal`, `Warning`, `Minor Problem`, `Major Problem`, and `Critical`.

Hardware Status

The server's hardware status. Possible values are: `Normal`, `Warning`, `Minor Problem`, `Major Problem`, and `Critical`.

Communication Status

The server's communication status. Possible values are: `Normal`, `Warning`, `Minor Problem`, `Major Problem`, and `Critical`.

Using the Basic Server Information window.

Most HPSS servers allow the SSM user to view the server's states and statuses.

Under *normal* conditions, the server states and statuses reported on the *Basic Server Information* window are as follows:

Operational State: `Enabled`

Usage State: `Active`

Administrative State: `Unlocked`

Execution State: `Active`

Service Status: `Normal`

Security Status: `Normal`

Software Status: Normal

Hardware Status: Normal

Communication Status: Normal

However, when the server is experiencing errors or encountering abnormal conditions, it will change the appropriate states and statuses to error values, notify SSM of the changes, and issue an alarm to SSM. Refer to *Section 9.4.2, “Alarm/Event Information”* for more information.



The Startup Daemon and the System Manager do not have their own Basic Server Information windows.

5.3.2. Specific server information

Using the server-specific information windows.

The majority of the HPSS servers also allow the SSM user to view and change their server-specific data through the SSM windows. A typical server allows authorized users to change the value of the fields and use them immediately in the current execution.

There are server-specific information windows for the following servers:

- Core Server
- Gatekeeper
- Location Server
- Migration/Purge Server
- Mover
- Physical Volume Library
- Physical Volume Repository
- RAIT Engine



*Any changes made on a server’s information window will be effective only during the server’s current execution. To make the changes permanent, make the identical changes on the server’s configuration window and if the setting changed allows, the change should be applied by clicking the **Apply Config** button on the HPSS Health and Status screen. Refer to Table 5.3, “Server reinitialization behavior” to determine what settings can be applied by **Apply Config** (i.e reinitialization). Otherwise, restart the server (see Section 5.2, “Server configuration” and below).*

If the server’s operational state is Broken or Unknown, the specific server information window will not be available for viewing.

The server-specific information windows are described in detail below.

5.3.2.1. Core Server Information window

The screenshot shows the 'Core Server Information' window with the following fields and values:

Core Server	Core Server
Global Database Name	cfgx
Subsystem Database Name	subsysx1
Schema Name	HPSS
Root Fileset Name	FilesetRoot.449
Root Fileset ID	3272831213, ,3838938124
Maximum Open Bitfiles	2000
Maximum Active I/O Reqs	190
Maximum Active Copy Reqs	95
COS Change Stream Count	1
COS Change Retry Limit	0
Tape Dismount Delay (seconds)	30
Tape Handoff Delay (seconds)	120
Tape Queue Method	Offset
Tape Queue Active Set Size	1024
Tape Max Reader Wait (seconds)	300
PVL Max Connection Wait (seconds)	300
Fragment Smallest Block	8 (2^3)
<input type="checkbox"/> Log ENOSPACE Details (debug only)	
Root User ID	0
<input checked="" type="checkbox"/> Root is Superuser	
I/O Operations to Log (debug only)	
	Read Write Copy
Disk	<input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/>
Tape	<input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/>
Classes of IODs/IORs to Log (debug only)	
<input type="checkbox"/> SS Client IODs/IORs <input type="checkbox"/> Mover IODs/IORs <input type="checkbox"/> Rait Engine IODs/IORs (tape only)	
<input type="checkbox"/> Passive Replies <input type="checkbox"/> Mover Abort IODs <input type="checkbox"/> Mover IOCTLS	
<div> <div>Name Space Statistics</div> <div>Disk/Tape Statistics</div> <div>Trashcan Statistics</div> <div>Options</div> </div> <div> <div>Subsystem Statistics</div> <div> <div>Stages</div> <div>0</div> </div> <div> <div>Migrations</div> <div>0</div> </div> <div> <div>Purges</div> <div>0</div> </div> <div> <div>File Deletes</div> <div>0</div> </div> <div> <div>Last Reset Time</div> <div>Jan 26, 2015 4:53:04 PM</div> </div> </div>	
<input type="checkbox"/> Freeze <input type="button" value="Update"/> <input type="button" value="Refresh"/> <input type="button" value="Dismiss"/>	
Retrieving data - Succeeded	

Field descriptions.

Core Server

The descriptive name of the Core Server.

Global Database Name

The name of the global database for the HPSS system.

Subsystem Database Name

The name of the database which contains the subsystem tables used by this Core Server.

Schema Name

The name of the database schema.

Root Fileset Name

The name of the root fileset used by the Core Server.

Root Fileset ID

The fileset id for the root fileset used by the Core Server.

Maximum Open Bitfiles

The maximum number of bitfiles that can be open simultaneously.

Maximum Active I/O Reqs

The maximum number of simultaneous I/O requests allowed. (See *Section 5.2.2.2, “Core Server I/O limit configuration”*. for a detailed description of how this field interacts with other configuration fields)

Maximum Active Copy Reqs

The maximum number of simultaneous copy requests allowed. (See *Section 5.2.2.2, “Core Server I/O limit configuration”*. for a detailed description of how this field interacts with other configuration fields)

COS Change Stream Count

Indicates the number of background COS change threads to run in the Core Server. The default is "1".

COS Change Retry Limit

Limits the number of times that the Core Server will retry a failed COS change operation. When this limit is reached, the attempt to change the COS of the file is terminated.

Tape Dismount Delay (seconds)

The amount of time, in seconds, a mounted tape volume will remain idle before being dismounted by the Core Server. Larger values may reduce undesirable tape dismount/remount events, at the expense of lower tape drive utilization.

Tape Handoff Delay (seconds)

The amount of time, in seconds, a mounted tape will be held in a client's session before become eligible to be handed off to another client that wishes to use the tape.

Tape Queue Method

This field may be set to either `FIFO` or `Offset`. If set to `FIFO`, the Core Server will process tape requests per VV in a FIFO fashion without doing any intelligent ordering. This somewhat mimics legacy behavior and can be useful for sites which want to closely manage their own ordering. If set to `Offset`, the Core Server will attempt to reorganize the VV session queue in order to improve performance. If available, this may use drive technologies such as Recommended Access Ordering (RAO).

If a file called `nodrive` exists in the HPSS tmp folder (`HPSS_PATH_TMP`), its presence will disable RAO-based sorting of tape requests and instead use strict offset ordering. If a file called `noschedule` exists in this folder it will disable all scheduling, resulting in a FIFO ordering for tape requests from that point until the file is removed. Note that this will not impact any existing schedules.

Individual VVs may also be toggled to FIFO scheduling using the **FIFO** check box on the VV information window. See the *Section 8.5.4.2, “Core Server Tape Volume information window”* for more information.

Tape Queue Active Set Size

Number of sessions which may be scheduled together. This yields a compromise between scheduling efficiency and potential wait times since this size is the maximum number of items within that set that may be scheduled ahead of any other request in that set. This may also be limited by the hardware limitations of Recommended Access Ordering (RAO) if the drive a tape is mounted on supports RAO.

Tape Max Reader Wait (seconds)

This is the amount of time, in seconds, that readers of a tape may spend waiting on writers before the scheduling priority shifts to prefer readers over writers for some period of time. There is a minimum of five minutes between reader/writer priority changes. Setting this to zero disables preference switching and will cause writers to always be preferred.

PVL Max Connection Wait (seconds)

The amount of time, in seconds, the Core Server will wait to connect to the PVL before declaring an error in pending PVL jobs.

Fragment Smallest Block

Fragment Smallest Block sets the boundary used to determine where to remove excess clusters from allocations that have excess space. The smallest block returned to the free space map from the excess space at the end of an extent will be this size or larger.

Log ENOSPACE Details (debug only)

If ON, the Core Server logs detailed debug class log messages about disk and tape allocation requests that end in `HPSS_ENOSPACE`. In normal operation, these details are too voluminous to log economically, so they are not logged by default. But if the cause of an `ENOSPACE` condition is not evident, setting this switch will enable logging of all of the details. Because of the volume of information logged, leaving this switch on for long periods of time is not recommended.

Root User ID

The numerical ID for the root user used by the Core Server.

Root is Superuser

If this flag is ON, clients using the root user ID will have superuser powers and permissions.

I/O Operations to Log (debug only) , Classes of IODs/IORs to Log (debug only)

These switches control optional logging in the Core Server. The Core Server always logs IODs and IORs when I/O operations fail, but does not log IODs and IORs for successful operations. These switches allow the administrator to turn on logging of IODs and IORs in successful tape and disk I/O operations for diagnostic purposes.

Normally these switches should be off to avoid logging large amounts of information that won't be useful. Changes to these switch settings in the managed object take effect immediately. Reinitialization of the server is not necessary.

Subsystem Statistics tab

Stages

The number of stage requests processed in the Core Server since startup or last reset of the statistics.

Migrations

The number of migrate requests processed in the Core Server since startup or last reset of the statistics.

Purges

The number of purge requests processed in the Core Server since startup or last reset of the statistics.

File Deletes

The number of bitfile delete requests processed in the Core Server since startup or last reset of the statistics.

Last Reset Time

The last time the subsystem statistics were reset. If this value is "0", the statistics have not been reset since server startup.

Core Server Information

File Edit Help

Core Server: Core Server

Global Database Name: cfgx

Subsystem Database Name: subsysx1

Schema Name: HPSS

Root Fileset Name: FilesetRoot.449

Root Fileset ID: 3272831213, ,3838938124

Maximum Open Bitfiles: 2000

Maximum Active I/O Reqs: 190

Maximum Active Copy Reqs: 95

COS Change Stream Count: 1

COS Change Retry Limit: 0

Tape Dismount Delay (seconds): 30

Tape Handoff Delay (seconds): 120

Tape Queue Method: Offset

Tape Queue Active Set Size: 1024

Tape Max Reader Wait (seconds): 300

PVL Max Connection Wait (seconds): 300

Fragment Smallest Block: 8 (2^3)

☐ Log ENOSPACE Details (debug only)

Root User ID: 0 ☒ Root is Superuser

I/O Operations to Log (debug only)

	Read	Write	Copy
Disk	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Tape	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Classes of IODs/IORs to Log (debug only)

☐ SS Client IODs/IORs ☐ Mover IODs/IORs ☐ Rait Engine IODs/IORs (tape only)

☐ Passive Replies ☐ Mover Abort IODs ☐ Mover IOCTLS

Name Space Statistics | **Disk/Tape Statistics** | **Trashcan Statistics** | **Options**

Subsystem Statistics

Files	33	Hard Links	0
Directories	23	Junctions	13
Symbolic Links	1	Filesets	2
		Total Name Space Objects	72

☐ Freeze

Retrieving data - Succeeded

Name Space Statistics tab

Files

The number of files managed by this Core Server.

Directories

The number of directories managed by this Core Server.

Symbolic Links

The number of symbolic links managed by this Core Server.

Hard Links

The number of hard links managed by this Core Server.

Junctions

The number of junctions managed by this Core Server.

Filesets

The number of filesets managed by this Core Server.

Total Name Space Objects

The total number of name space records in the database. Note that this total may not match the total achieved by summing the number of the above objects. Some objects, such as hard links, require more than one record.

Core Server Information

File Edit Help

Core Server

Core Server

Global Database Name

cfg

Subsystem Database Name

subsys1

Schema Name

HPSS

Root Fileset Name

FilesetRoot.3539

Root Fileset ID

3272984902,,2667178584

Maximum Open Bitfiles

200000

Maximum Active I/O Reqs

80

Maximum Active Copy Reqs

80

COS Change Stream Count

1

COS Change Retry Limit

0

Tape Dismount Delay (seconds)

600

Tape Handoff Delay (seconds)

300

Tape Queue Method

Offset

Tape Queue Active Set Size

1024

Tape Max Reader Wait (seconds)

300

PVL Max Connection Wait (seconds)

300

Fragment Smallest Block

1 (2^0)

☐ Log ENOSPACE Details (debug only)

Root User ID

0

☒ Root is Superuser

I/O Operations to Log (debug only)

Read

Write

Copy

Disk

☐

☒

☒

Tape

☒

☐

☒

Classes of IODs/IORs to Log (debug only)

☒ SS Client IODs/IORs

☒ Mover IODs/IORs

☒ Rait Engine IODs/IORs (tape only)

☒ Passive Replies

☐ Mover Abort IODs

☐ Mover IOCTLS

Subsystem Statistics

Name Space Statistics

Disk/Tape Statistics

Trashcan Statistics

Options

Total Disk Virtual Volumes

3

Active Disk Bytes

62,019,076,096

Total Tape Virtual Volumes

38

Free Disk Bytes

2,137,105,891,328

Tape Aggregates

1,881

Active Tape Bytes

15,347,628,285

File Segments in Tape Aggregates

2,034,598

Free Tape Bytes

44,667,659,878,400

Bytes in Tape Aggregates

4,301,097,060

☐ Freeze

Update

Refresh

Dismiss

Retrieving data - Succeeded

Disk/Tape Statistics tab

Total Disk Virtual Volumes

The number of disk virtual volumes managed by this Core Server.

Total Tape Virtual Volumes

The number of tape virtual volumes managed by this Core Server.

Tape Aggregates

The number of tape aggregates managed by this Core Server. A tape aggregate is a single tape storage segment that contains more than one file.

File Segments in Tape Aggregates

The number of file segments in tape aggregates managed by this Core Server. This is a count of files within all aggregates managed by the server. Files may be split across tape aggregates so may be counted multiple times in this statistic.

Bytes in Tape Aggregates

The number of bytes stored in the tape aggregates managed by this Core Server.

Active Disk Bytes

This is the sum of the bytes in the storage segments assigned to files, including slack space. Slack space is the disk space assigned to disk storage segments that is not occupied with file data. The sum of the lengths of files on disk will be less than the **Active Disk Bytes** because of the presence of slack space.

Free Disk Bytes

This is the sum of the bytes not assigned to files. More precisely, it is the sum of the disk volume bytes not assigned to disk storage segments.

Active Tape Bytes

This is the sum of the bytes assigned to files on tapes. It does not include deleted file space that has not been reclaimed.

Free Tape Bytes

This is an estimate based on the sum of the estimated sizes of the partially written and unwritten tape volumes. It is not, and cannot be, an accurate value as the amount of data that can be written on tapes varies with individual tape volumes and data compression levels.

Core Server Information

File Edit Help

Core Server: Core Server

Global Database Name: cfgx

Subsystem Database Name: subsysx1

Schema Name: HPSS

Root Fileset Name: FilesetRoot.449

Root Fileset ID: 3272831213, ,3838938124

Maximum Open Bitfiles: 2000

Maximum Active I/O Reqs: 190

Maximum Active Copy Reqs: 95

COS Change Stream Count: 1

COS Change Retry Limit: 0

Tape Dismount Delay (seconds): 30

Tape Handoff Delay (seconds): 120

Tape Queue Method: Offset

Tape Queue Active Set Size: 1024

Tape Max Reader Wait (seconds): 300

PVL Max Connection Wait (seconds): 300

Fragment Smallest Block: 8 (2^3)

☐ Log ENOSPACE Details (debug only)

Root User ID: 0 ☒ Root is Superuser

I/O Operations to Log (debug only)

	Read	Write	Copy
Disk	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Tape	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Classes of IODs/IORs to Log (debug only)

☐ SS Client IODs/IORs ☐ Mover IODs/IORs ☐ Rait Engine IODs/IORs (tape only)

☐ Passive Replies ☐ Mover Abort IODs ☐ Mover IOCTLS

Name Space Statistics **Disk/Tape Statistics** **Trashcan Statistics** **Options**

Subsystem Statistics

Last Trash Statistics Time	Jan 27, 2015 11:45:19 AM
Total Objects in the Trash	0
Size of Objects in the Trash	0
Total Objects Eligible for Unlink	0
Size of Objects Eligible for Unlink	0

☐ Freeze

Retrieving data - Succeeded

Trashcan Statistics tab

Last Trash Statistics Time

The last time that the background thread ran which periodically gathers trashcan-related statistics.

Total Objects in the Trash

This is a count of the total number of entries currently in all trashcans (the NSTRASH database table).

Size of Objects in the Trash

The total number of bytes used by all files currently in trashcans (the NSTRASH database table).

Total Objects Eligible for Unlink

The count of the number of entries in the NSTRASH table that are currently eligible for permanent deletion from trashcans (the NSTRASH table). Entries in trashcans become eligible for deletion when they have been in a trashcan for **Trash Unlink Eligible Time** seconds.

Trash Unlink Eligible Time is a field on the *Global Configuration* window. (Refer to *Section 4.1, “Global Configuration window”*.)

Size of Objects Eligible for Unlink

The total number of bytes used by the trashcan file objects that are currently eligible for permanent deletion.

Core Server Information

File Edit Help

Core Server: Core Server

Global Database Name: cfgx

Subsystem Database Name: subsysx1

Schema Name: HPSS

Root Fileset Name: FilesetRoot.449

Root Fileset ID: 3272831213, ,3838938124

Maximum Open Bitfiles: 2000

Maximum Active I/O Reqs: 190

Maximum Active Copy Reqs: 95

COS Change Stream Count: 1

COS Change Retry Limit: 0

Tape Dismount Delay (seconds): 30

Tape Handoff Delay (seconds): 120

Tape Queue Method: Offset

Tape Queue Active Set Size: 1024

Tape Max Reader Wait (seconds): 300

PVL Max Connection Wait (seconds): 300

Fragment Smallest Block: 8 (2^3)

☐ Log ENOSPACE Details (debug only)

Root User ID: 0 ☒ Root is Superuser

I/O Operations to Log (debug only)

	Read	Write	Copy
Disk	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Tape	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Classes of IODs/IORs to Log (debug only)

☐ SS Client IODs/IORs ☐ Mover IODs/IORs ☐ Rait Engine IODs/IORs (tape only)

☐ Passive Replies ☐ Mover Abort IODs ☐ Mover IOCTLS

Name Space Statistics **Disk/Tape Statistics** **Trashcan Statistics** **Options**

Subsystem Statistics

☐ Can change UID to self if has Control Perm

☐ Can change UID if has Delete Perm on Security ACL

☒ Object names can contain unprintable characters

☐ COS Copy to Disk

☐ Account Flag Switch

☐ Freeze

Retrieving data - Succeeded

Options tab

Can change UID to self if has Control Perm

If this flag is ON, any user having Control permission to an object can change the UID of that object to their own (but not any other) UID. This field can be changed only from the *Global Configuration* window. It is displayed here for informational purposes only.

Can change UID if has Delete Perm on Security ACL

If this flag is ON, any principal found on the Core Server Security ACL having Delete (and only Delete) permission may change the UID of an object they own to any other UID. This field can be changed only from the *Global Configuration* window. It is displayed here for informational purposes only.

Object names can contain unprintable characters

If this flag is ON, no check will be made to ensure that new object and fileset names contain only printable characters. If this flag is OFF, new object and fileset names must be composed of only printable characters and if any unprintable characters are detected in these new names, an error will be returned. Printable characters are defined as all characters in the range from 0x20 to 0x7E (inclusive) in the 7-bit ASCII character set. This field can be changed only from the *Global Configuration* window. It is displayed here for informational purposes only.

COS Copy to Disk

If ON, all copy operations associated with COS changes should be directed to disk if the hierarchy of the target COS has a level 0 disk storage class. This flag may be changed only from the Specific tab of the *Core Server Configuration* window. It is displayed here for informational purposes only.

Account Flag Switch

If this flag is ON, it indicates to use the second accounting bit in the bitfile metadata table to indicate that the bitfile has valid accounting metadata. If the flag is OFF, it indicates to use the first bit. The default is OFF. This flag setting is changeable only via special procedures. It is provided to allow a site to recover or rebuild its accounting metadata if this becomes necessary. Special procedures are required for this. Contact HPSS support to do this.

5.3.2.2. Gatekeeper Information window

The screenshot shows a window titled "Gatekeeper Information" with a menu bar (File, Edit, Help). It contains several input fields and a table of statistics.

Server Name: Gatekeeper

Default Wait Time (seconds): 10

Site Policy Pathname (UNIX): /var/hpss/gk/gksitepolicy

Administrative Activity Statistics

	Calls	Errors	Retries
Get Basic Server Info	0	0	
Set Basic Server Info	2	0	
Get Gatekeeper Server Info	0	0	
Set Gatekeeper Server Info	4	0	
Creates	0	0	0
Auth Caller Creates	0	0	
Creates Completed	0	0	
Opens	0	0	0
Auth Caller Opens	0	0	
Closes	0	0	
Stages	0	0	0
Auth Caller Stages	0	0	
Stages Completed	0	0	
Client Clean Ups	0	0	
Get Monitor Types	0	0	
Pass Thrus	0	0	
Queries	0	0	
Read Site Policies	0	0	

Jan 28, 2013 2:48:11 PM Last Reset Time

Reset

Read Site Policy

☐ Freeze Update Refresh Dismiss

Retrieving data - Succeeded

Field descriptions.

Server Name

The descriptive name of the Gatekeeper.

Default Wait Time (seconds)

The number of seconds the client must wait before retrying a request. This value must be greater than zero and is used if the Gatekeeping Site Interface returns a wait time of zero for the create, open, or stage request being retried. Changing the value of this field will cause the Gatekeeper to use the new value until the next restart at which point it will then go back to using the value defined in the *Gatekeeper Configuration* window. Refer to Section 5.2.3, "Gatekeeper-specific configuration".

Site Policy Pathname (UNIX)

This field can only be set from the *Gatekeeper Configuration* window. Refer to Section 5.2.3, "Gatekeeper-specific configuration".

Administrative activity statistics

These fields are internal volatile statistics for the Gatekeeper. Each row describes a different API. There are three columns: Calls, Errors and Retries. Each of these columns represent a count of the number of calls, errors, and retries since the statistics were reset. The Calls column is the number of times the API was called. The Errors column is the number of times the API call resulted in an error being returned. The Retries column is the number of times that the API call resulted in the HPSS_ERETRY error being returned to the Client API. The Retries column only applies to API requests that may return the HPSS_ERETRY status. A retry is not counted as an error.

The statistics are reset to zero whenever the Gatekeeper is recycled or the **Reset** button is clicked.

Get Basic Server Info

Statistics from the gk_admin_ServerGetAttrs API.

Set Basic Server Info

Statistics from the gk_admin_ServerSetAttrs API. This API is called by the SSM System Manager when the Gatekeeper's *Basic Server Information* window is opened.

Get Gatekeeper Server Info

Statistics from the gk_admin_GKGetAttrs API.

Set Gatekeeper Server Info

Statistics from the gk_admin_GKSetAttrs API. This API is called by the SSM System Manager when the *Gatekeeper Information* window is opened.

Creates

Statistics from the gk_Create API. This API is called by the Core Server when the Gatekeeper is monitoring Create Requests.

Auth Caller Creates

Statistics from authorized caller (for example, MPS) calls to the gk_Create API.

Creates Completed

Statistics from the gk_CreateComplete API. This API is called by the Core Server when the Gatekeeper is monitoring Create Requests and the create completes.

Opens

Statistics from the gk_Open API. This API is called by the Core Server when the Gatekeeper is monitoring Open Requests.

Auth Caller Opens

Statistics from authorized caller (for example, MPS) calls to the gk_Open API.

Closes

Statistics from the gk_Close API. This API is called by the Core Server when the Gatekeeper is monitoring Open Requests and the file is closed.

Stages

Statistics from the gk_Stage API. This API is called by the Core Server when the Gatekeeper is monitoring Stage Requests.

Auth Caller Stages

Statistics from authorized caller (for example, MPS) calls to the gk_Stage API.

Stages Completed

Statistics from the gk_StageComplete API. This API is called by the Core Server when the Gatekeeper is monitoring Stage Requests and the stage completes.

Client Clean Ups

Statistics from the gk_CleanUpClient API. This API is called by the Core Server when the Gatekeeper is monitoring Requests and a client disconnects.

Get Monitor Types

Statistics from the gk_GetMonitorTypes API. This API is called by the Core Server to figure out what types of Requests being monitored by the Gatekeeper.

Pass Thrus

Statistics from the gk_PassThru API.

Queries

Statistics from the gk_Query API.

Read Site Policies

Statistics from the gk_ReadSitePolicy API.

Last Reset Time

The time stamp when the statistics were last reset to zero.

Associated button descriptions.**Reset**

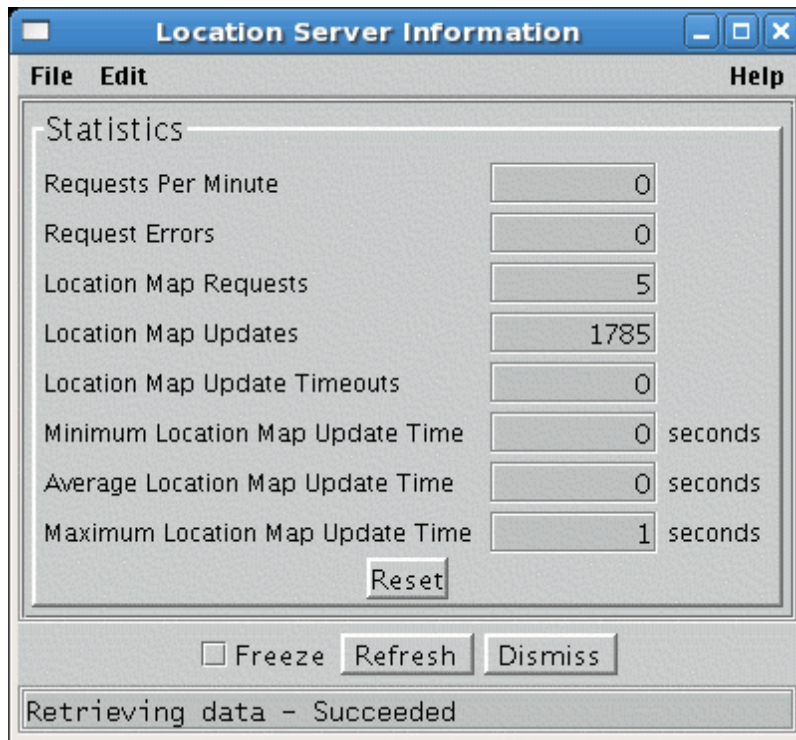
Reset the administrative activity statistics to "0" and the **Last Reset Time** to the current date and time.

Read Site Policy

This button is used to inform the GatekeepingSite Interface (gk_site_ReadSitePolicy) of a policy change to the **Site Policy Pathname** file.

Note: a "policy change" cannot include a change in the types of requests being monitored. When changing the types of requests being monitored (authorized caller, create, open, and stage), the Gatekeeper must be recycled so that the Core Server can learn which request types need to include Gatekeeping services.

5.3.2.3. Location Server Information window



The *Location Server Information* window displays statistics for the server's request processing and background updating functions.

The maximum number of requests a single Location Server can process in one minute varies depending on the machine load caused by other processes as well as the CPU performance. Generally the Location Server will report warning messages when it becomes loaded.

If timeouts occur, consider increasing the **Location Map Timeout** field on the *Location Policy* window.

Field descriptions.

Requests Per Minute

The number of client requests received during the last minute.

Request Errors

The number of errors returned to client requests.

Location Map Requests

The number of client requests received for local server location information (maps).

Location Map Updates

The number of background location map updates that have occurred.

Location Map Update Timeouts

The number of timeouts that have occurred while waiting for location map information.

Minimum Location Map Update Time

The shortest time, in seconds, needed for a location map update.

Average Location Map Update Time

The average time, in seconds, needed for a location map update.

Maximum Location Map Update Time

The longest time, in seconds, needed for a location map update.

Associated button descriptions.**Reset**

Reset the statistics to "0".

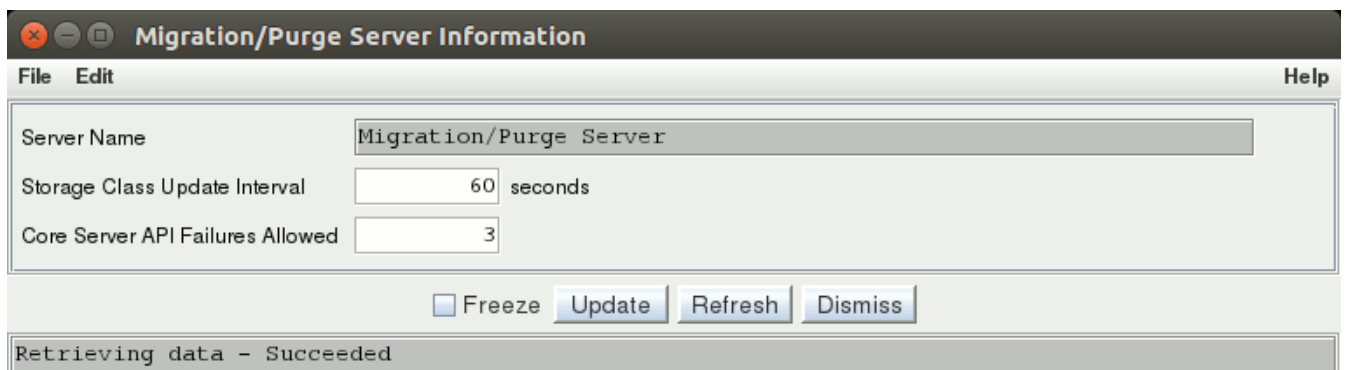
Related information.

HPSS Error Manual : Problem diagnosis and resolution chapter : Location Server problems section.

5.3.2.4. Migration/Purge Server Information window

This window allows you to view and update the type-specific information associated with a Migration/Purge Server (MPS).

Any changes made to fields on this window are sent directly to the MPS and are effective immediately.

**Field descriptions.****Server Name**

The descriptive name of the MPS.

Storage Class Update Interval

The interval that indicates how often the MPS will query the Core Server in its subsystem to get the latest storage class statistics. This is also the interval the MPS uses to check whether it needs to initiate a purge operation on the storage class based on the associated purge policy. The valid range for this field is 10 - 600 seconds.

Core Server API Failures Allowed

The maximum number of consecutive Core Server API failures allowed before MPS aborts a migration or purge run. This applies to both disk and tape migration as well as disk purge. If this limit is reached during a disk migration run, migration skips to the next hierarchy to which that disk storage class belongs. Once the run has attempted to migrate all such hierarchies the run will

end. If this limit is reached during a disk purge or tape migration run, the migration or purge run aborts.

5.3.2.5. Mover Information window

This window allows you to view and update the type-specific information associated with a Mover. Any changes made to fields on this window are sent directly to the Mover after the appropriate button is clicked and are effective immediately.

Field descriptions.

Server Name

The descriptive name of the Mover.

Number of Request Tasks

The number of Mover request-processing tasks that currently exist.

Number of Active Requests

The number of active requests that are being handled by the Mover.

Time of Last Statistics Reset

The time and date when the Mover statistics were last reset.

Buffer Size

The size, in bytes, of each buffer used by the Mover when transferring data. Each Mover task uses two buffers of this size to perform double buffering during I/O operations. See the *Mover Configuration* window **Buffer Size** field described in *Section 5.2.6, “Mover-specific configuration”* for more detailed information.

Number of Requests Processed

The number of requests that the Mover has processed since the last time this field was cleared. Click on the **Reset** button to the right of the field to reset the value to zero.

Number of Request Errors

The number of requests that have returned errors since the last time this field was cleared. Click on the **Reset** button to the right of the field to reset the value to zero.

Number of Data Transfers

The number of data transfers that the Mover has handled since the last time this field was cleared. Click on the **Reset** button to the right of the field to reset the value to zero.

Number of Bytes Moved

The number of bytes of data that the Mover has transferred since the last time this field was cleared. Click on the **Reset** button to the right of the field to reset the value to zero.

Total Move Time

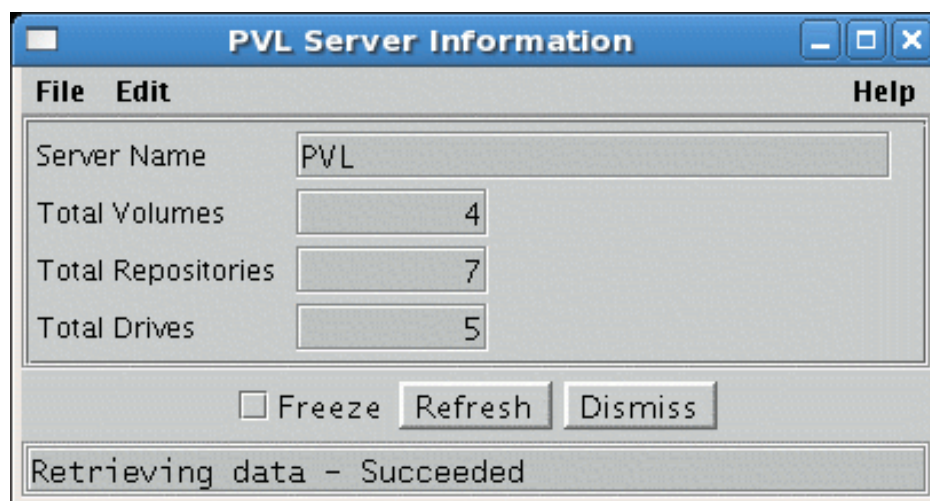
The total time, in days, hours, minutes, and seconds, that the Mover tasks have spent transferring data since the last time this field was cleared. Click on the **Reset** button to the right of the field to reset the value to zero.

Operational notes:

1. The Mover statistics (**Number of Requests Processed**, **Number of Request Errors**, **Number of Data Transfers**, **Number of Bytes Moved**, and **Total Move Time**) may all be reset to zero. If any of these fields are reset the **Time of Last Statistics Reset** field will also be updated with the current time. The statistics fields are all initialized to zero when the Mover is started.
2. The **Total Move Time** statistic should not be used to determine Mover throughput rates if the Mover handles striped media, since this is a total for all Mover tasks involved in the transfer. This total will most likely be considerably greater than the elapsed time necessary to transfer the data.

Related information.

HPSS Error Manual : Problem diagnosis and resolution chapter : Mover problems section.

5.3.2.6. Physical Volume Library (PVL) information window

This window allows you to view the type-specific information associated with a PVL.

Field descriptions.**Server Name**

The descriptive name of the PVL.

Total Volumes

The total number of volumes that have been imported into the PVL.

Total Repositories

The total number of PVRs in the *Servers* list window.

Total Drives

The total number of drives controlled by this PVL.

5.3.2.7. Operator PVR information window

Operator PVR Server Information

File Edit Help

Server Name Operator PVR

Total Cartridges 0

Cartridge Capacity 1000

Cartridge Alarm Threshold 90 percent

Shelf Tape Check-In Retry 30 seconds

Shelf Tape Check-In Alarm 10 minutes

Dismount Delay 15 minutes

Defer Dismount Exempt Count -1

Characteristics

☐ Defer Dismounts ☐ Support Shelf Tape

☐ Freeze Update Refresh Dismiss

Retrieving data - Succeeded

This window allows you to view and update the type-specific information associated with an Operator PVR. Any changes made to fields on this window are sent directly to the PVR and are effective immediately.

Field descriptions.**Server Name**

The descriptive name of the PVR.

Total Cartridges

The number of cartridges currently being managed by the PVR.

Cartridge Capacity

The total number of cartridge slots in the library dedicated to this HPSS PVR. This may or may not be the total cartridge capacity of the library; a site might use part of the library for some other HPSS PVR or for some non-HPSS application. The PVR uses the **Cartridge Capacity** field and the **Cartridge Alarm Threshold** field to determine when to send an alarm that the total cartridge threshold has been exceeded.

Cartridge Alarm Threshold

The percentage of the **Cartridge Capacity** at which the PVR will send an alarm.

Shelf Tape Check-In Retry

The number of seconds the PVR will wait before asking the robot if a requested shelf tape has been checked in. The PVR will continue checking at this interval until the tape is checked in. This field applies only if the **Support Shelf Tape** check box is selected. The retry value must be "30" or greater.

Shelf Tape Check-In Alarm

The PVR will periodically log alarm messages when a requested shelf tape has not been checked in. This field specifies the number of minutes between alarms. This field is only active if the **Support Shelf Tape** check box is selected. The alarm value must be "2" or greater.

Dismount Delay

When **Defer Dismounts** is checked, this value is used by the PVL to determine the number of minutes that dismounts are delayed after the last data access.

Defer Dismount Exempt Count

If set greater than "0", this value will specify the number of drives within a PVR that will be exempt from the defer dismount logic. This behavior will take place for all drives within the PVR with a drive/device "Library Unit ID" set (see Drive/Device Configuration). If multiple Library Unit IDs are set, the same Exempt count will apply to all Library Unit ID subsets.

The goal of this option is to allow a finer control of the defer dismount logic to account for transfer costs (for example, elevator, pass through costs). Defer dismount logic to apply logic on a sub library case (unit). The definition of a "unit" is placed in the hands of the administrator and likely based on library characteristics and topology.

Characteristics

Flags for the PVR.

Defer Dismounts

If ON, the PVL will delay the dismounting of a tape cartridge until the drive is required by another job or until the **Dismount Delay** time limit is exceeded. If the **Defer Dismount Exempt Count** is also set, the defer dismount logic will not be exercised on the number of drives set. This may result in free drives available for new mount requests without first generating defer dismount operations. See **Defer Dismount Exempt Count** for details.

A PVL restart is required in order for changes to this flag to take effect.

Support Shelf Tape

If ON, the PVR and the PVL will support the removal of cartridges from the tape library using the **shelf_tape** utility.

5.3.2.8. AML PVR information window

Field	Value	Unit
Server Name	AML PVR	
Total Cartridges	0	
Cartridge Capacity	3000	
Cartridge Alarm Threshold	90	percent
Same Job on Controller	10	
Other Job on Controller	5	
Distance To Drive	2	
Shelf Tape Check-In Retry	30	seconds
Shelf Tape Check-In Alarm	10	minutes
Dismount Delay	15	minutes
Retry Mount Time Limit	-1	minutes
Drive Error Limit	-1	
Defer Dismount Exempt Count	-1	

Characteristics

☐ Defer Dismounts ☐ Support Shelf Tape

Client Name: _____

Server Name: _____

☐ Freeze U Update R Refresh D Dismiss

Retrieving data - Succeeded

This window allows you to view and update the type-specific information associated with an AML PVR. Any changes made to fields on this window are sent directly to the PVR and are effective immediately.

Field descriptions.

Server Name. The descriptive name of the PVR.

Total Cartridges

The number of cartridges currently being managed by the PVR.

Cartridge Capacity

The total number of cartridge slots in the library dedicated to this HPSS PVR. This may or may not be the total cartridge capacity of the library; a site might use part of the library for some other

HPSS PVR or for some non-HPSS application. The PVR uses the **Cartridge Capacity** field and the **Cartridge Alarm Threshold** field to determine when to send an alarm that the total cartridge threshold has been exceeded.

Cartridge Alarm Threshold

The percentage of the **Cartridge Capacity** at which the PVR will send an alarm.

Same Job on Controller, Other Job on Controller, Distance To Drive

These three values are used by the PVR when selecting a drive for a tape mount operation. The three values are essentially weights that are used to compute an overall score for each possible drive. After the score has been calculated, the drive with the lowest score is selected for the mount. If two or more drives tie for the lowest score, one drive is selected at random. The score is calculated as follows:

$$\begin{aligned} \text{Score} = & \\ & \text{Weight 1} \times \text{Cartridges from this job mounted on this drive's controller} + \\ & \text{Weight 2} \times \text{Cartridges from other jobs mounted on this drive's controller} + \\ & \text{Weight 3} \times \text{Units of distance from the cartridge to the drive} \end{aligned}$$

This method has the effect of distributing a striped tape mount across as many controllers as possible for the best performance. It also will try to pick controllers that are currently driving a minimum number of tapes. So, in an environment with many tape drives per controller, the above algorithm will minimize the load on any one controller.

The **Distance To Drive** helps minimize mount times by mounting the tape in a physically close drive. All other things being equal, the tape will be mounted in the closest drive. However, if the closest drive is attached to a heavily used controller, then a more distant drive will be selected.

Shelf Tape Check-In Retry

The number of seconds the PVR will wait before asking the robot if a requested shelf tape has been checked in. The PVR will continue checking at this interval until the tape is checked in. This field applies only if the **Support Shelf Tape** check box is selected. The retry value must be "30" or greater.

Shelf Tape Check-In Alarm

The PVR will periodically log alarm messages when a requested shelf tape has not been checked in. This field specifies the number of minutes between alarms. This field is only active if the **Support Shelf Tape** check box is selected. The alarm value must be "2" or greater.

Dismount Delay

When **Defer Dismounts** is checked, this value is used by the PVL to determine the number of minutes that dismounts are delayed after the last data access.

Retry Mount Time Limit

The default value for this field is "-1". When the default value is used, if an error is encountered during a PVR mount operation, the mount will pend and be retried every five minutes. Setting a value in this field will change the mount behavior to periodically retry the mount until the specified time limit is exceeded. Once exceeded, an error is generated and the mount request is canceled. If the mount request would have resulted in a write operation, the error returned will cause the Core Server to set the **VV Condition** of the associated tape volume to DOWN. Once

in DOWN state, the volume will no longer be available for read or write operations. For further information about the Core Server VV Condition, see *Section 8.5.4.2, “Core Server Tape Volume information window”*.

Drive Error Limit

This field is used in conjunction with the PVR Server **Retry Mount Time Limit**. If the number of consecutive mount errors which occur to any drive in this PVR equal or exceed this value, the drive is automatically locked by the PVL. The only mount errors that apply are those set through the **Retry Mount Time Limit** mechanism. The **Drive Error Count** field in the PVL Drive Information records the number of consecutive errors on a drive-by-drive basis. To turn off the automatic drive disable feature, set the **Drive Error Limit** to "0" or "-1". Changing this value will not change drive disable behavior until the PVL is recycled.

Defer Dismount Exempt Count

If set greater than "0", this value will specify the number of drives within a PVR that will be exempt from the defer dismount logic. This behavior will take place for all drives within the PVR with a drive/device "Library Unit ID" set (see Drive/Device Configuration). If multiple Library Unit IDs are set, the same Exempt count will apply to all Library Unit ID subsets.

The goal of this option is to allow a finer control of the defer dismount logic to account for transfer costs (for example, elevator, pass through costs). Defer dismount logic to apply logic on a sub library case (unit). The definition of a "unit" is placed in the hands of the administrator and likely based on library characteristics and topology.

Characteristics

Flags for the PVR.

Defer Dismounts

If ON, the PVL will delay the dismounting of a tape cartridge until the drive is required by another job or until the **Dismount Delay** time limit is exceeded. If the **Defer Dismount Exempt Count** is also set, the defer dismount logic will not be exercised on the number of drives set. This may result in free drives available for new mount requests without first generating defer dismount operations. See **Defer Dismount Exempt Count** for details.

A PVL restart is required in order for changes to this flag to take effect.

Support Shelf Tape

If ON, the PVR and the PVL will support the removal of cartridges from the tape library using the **shelf_tape** utility.

Client Name

The name of the client requesting authorization from the Distributed Automated Media Library Server.

Server Name

TCP/IP hostname or IP address of the AML OS/2-PC DAS server.

5.3.2.9. SCSI PVR information window

SCSI PVR Server Information

File Edit Help

Server Name	SCSI PVR
Total Cartridges	0
Cartridge Capacity	440
Cartridge Alarm Threshold	90 percent
Same Job on Controller	10
Other Job on Controller	5
Distance To Drive	0
Shelf Tape Check-In Retry	30 seconds
Shelf Tape Check-In Alarm	10 minutes
Dismount Delay	15 minutes
Retry Mount Time Limit	-1 minutes
Drive Error Limit	-1
Defer Dismount Exempt Count	-1

Characteristics

☐ Defer Dismounts ☐ Support Shelf Tape ☐ Enforce Home Location

Serial Number 0	/dev/fscsi2/d6c1c007100000000000
Serial Number 1	
Serial Number 2	
Serial Number 3	
Serial Number 4	
Serial Number 5	
Serial Number 6	
Serial Number 7	

☐ Freeze

Retrieving data - Succeeded

This window allows you to view and update the type-specific information associated with a SCSI PVR. Any changes made to fields on this window are sent directly to the PVR and are effective immediately.

Field descriptions.

Server Name

The descriptive name of the PVR.

Total Cartridges

The number of cartridges currently being managed by the PVR.

Cartridge Capacity

The total number of cartridge slots in the library dedicated to this HPSS PVR. This may or may not be the total cartridge capacity of the library; a site might use part of the library for some other HPSS PVR or for some non-HPSS application. The PVR uses the **Cartridge Capacity** field and the **Cartridge Alarm Threshold** field to determine when to send an alarm that the total cartridge threshold has been exceeded.

Cartridge Alarm Threshold

The percentage of the **Cartridge Capacity** at which the PVR will send an alarm.

Same Job on Controller, Other Job on Controller, Distance To Drive

These three values are used by the PVR when selecting a drive for a tape mount operation. The three values are essentially weights that are used to compute an overall score for each possible drive. After the score has been calculated, the drive with the lowest score is selected for the mount. If two or more drives tie for the lowest score, one drive is selected at random. The score is calculated as follows:

$$\begin{aligned} \text{Score} = & \\ & \text{Weight 1} \times \text{Cartridges from this job mounted on this drive's controller} + \\ & \text{Weight 2} \times \text{Cartridges from other jobs mounted on this drive's controller} + \\ & \text{Weight 3} \times \text{Units of distance from the cartridge to the drive} \end{aligned}$$

This method has the effect of distributing a striped tape mount across as many controllers as possible for the best performance. It also will try to pick controllers that are currently driving a minimum number of tapes. So, in an environment with many tape drives per controller, the above algorithm will minimize the load on any one controller.

The **Distance To Drive** helps minimize mount times by mounting the tape in a physically close drive. All other things being equal, the tape will be mounted in the closest drive. However, if the closest drive is attached to a heavily used controller, then a more distant drive will be selected.

Shelf Tape Check-In Retry

The number of seconds the PVR will wait before asking the robot if a requested shelf tape has been checked in. The PVR will continue checking at this interval until the tape is checked in. This field applies only if the **Support Shelf Tape** check box is selected. The retry value must be "30" or greater.

Shelf Tape Check-In Alarm

The PVR will periodically log alarm messages when a requested shelf tape has not been checked in. This field specifies the number of minutes between alarms. This field is only active if the **Support Shelf Tape** check box is selected. The alarm value must be "2" or greater.

Dismount Delay

When **Defer Dismounts** is checked, this value is used by the PVL to determine the number of minutes that dismounts are delayed after the last data access.

Retry Mount Time Limit

The default value for this field is "-1". When the default value is used, if an error is encountered during a PVR mount operation, the mount will pend and be retried every five minutes. Setting

a value in this field will change the mount behavior to periodically retry the mount until the specified time limit is exceeded. Once exceeded, an error is generated and the mount request is canceled. If the mount request would have resulted in a write operation, the error returned will cause the Core Server to set the **VV Condition** of the associated tape volume to DOWN. Once in DOWN state, the volume will no longer be available for read or write operations. For further information about the Core Server VV Condition, see *Section 8.5.4.2, "Core Server Tape Volume information window"*.

Drive Error Limit

This field is used in conjunction with the PVR Server **Retry Mount Time Limit**. If the number of consecutive mount errors which occur to any drive in this PVR equal or exceed this value, the drive is automatically locked by the PVL. The only mount errors that apply are those set through the **Retry Mount Time Limit** mechanism. The **Drive Error Count** field in the PVL Drive Information records the number of consecutive errors on a drive by drive basis. To turn off the automatic drive disable feature, set the **Drive Error Limit** to "0" or "-1". Changing its value will not change drive disable behavior until the PVL is recycled.

Defer Dismount Exempt Count

If set greater than "0", this value will specify the number of drives within a PVR that will be exempt from the defer dismount logic. This behavior will take place for all drives within the PVR with a drive/device "Library Unit ID" set (see Drive/Device Configuration). If multiple Library Unit IDs are set, the same Exempt count will apply to all Library Unit ID subsets.

The goal of this option is to allow a finer control of the defer dismount logic to account for transfer costs (for example, elevator, pass through costs). Defer dismount logic to apply logic on a sub library case (unit). The definition of a "unit" is placed in the hands of the administrator and likely based on library characteristics and topology.

Characteristics

Flags for the PVR.

Defer Dismounts

If ON, the PVL will delay the dismounting of a tape cartridge until the drive is required by another job or until the **Dismount Delay** time limit is exceeded. If the **Defer Dismount Exempt Count** is also set, the defer dismount logic will not be exercised on the number of drives set. This may result in free drives available for new mount requests without first generating defer dismount operations. See **Defer Dismount Exempt Count** for details.

A PVL restart is required in order for changes to this flag to take effect.

Support Shelf Tape

If ON, the PVR and the PVL will support the removal of cartridges from the tape library using the **shelf_tape** utility.

Enforce Home Location

If ON, the SCSI PVR will always try to dismount a mounted cart back to its home location. Otherwise, it will just use the first free slot. The **scsi_home** utility can be used to view the home location values.

Serial Number 0 through Serial Number 7

The serial numbers of the robots the SCSI PVR is controlling. Used to detect control paths to the robots. Only Serial Number 0 is valid at this time.

5.3.2.10. STK PVR information window

STK PVR Server Information

File Edit Help

Server Name: STK PVR

Total Cartridges: 0

Cartridge Capacity: 6000

Cartridge Alarm Threshold: 90 percent

Same Job on Controller: 10

Other Job on Controller: 5

Distance To Drive: 20

Shelf Tape Check-In Retry: 30 seconds

Shelf Tape Check-In Alarm: 10 minutes

Dismount Delay: 15 minutes

Retry Mount Time Limit: -1 minutes

Drive Error Limit: -1

Defer Dismount Exempt Count: -1

Characteristics

☐ Defer Dismounts ☐ Support Shelf Tape

ACSL5 Packet Version: 4

Alternate SSI Port:

☐ Freeze Update Refresh Dismiss

Retrieving data - Succeeded

This window allows you to view and/update the type-specific information associated with a STK PVR. Any changes made to fields on this window are sent directly to the PVR and are effective immediately.

Field descriptions.

Server Name

The descriptive name of the PVR.

Total Cartridges

The number of cartridges currently being managed by the PVR.

Cartridge Capacity

The total number of cartridge slots in the library dedicated to this HPSS PVR. This may or may not be the total cartridge capacity of the library; a site might use part of the library for some other HPSS PVR or for some non-HPSS application. The PVR uses the **Cartridge Capacity** field and the **Cartridge Alarm Threshold** field to determine when to send an alarm that the total cartridge threshold has been exceeded.

Cartridge Alarm Threshold

The percentage of the **Cartridge Capacity** at which the PVR will send an alarm.

Same Job on Controller, Other Job on Controller, Distance To Drive

These three values are used by the PVR when selecting a drive for a tape mount operation. The three values are essentially weights that are used to compute an overall score for each possible drive. After the score has been calculated, the drive with the lowest score is selected for the mount. If two or more drives tie for the lowest score, one drive is selected at random. The score is calculated as follows:

$$\begin{aligned} \text{Score} = & \\ & \text{Weight 1} \times \text{Cartridges from this job mounted on this drive's controller} + \\ & \text{Weight 2} \times \text{Cartridges from other jobs mounted on this drive's controller} + \\ & \text{Weight 3} \times \text{Units of distance from the cartridge to the drive} \end{aligned}$$

This method has the effect of distributing a striped tape mount across as many controllers as possible for the best performance. It also will try to pick controllers that are currently driving a minimum number of tapes. So, in an environment with many tape drives per controller, the above algorithm will minimize the load on any one controller.

The **Distance To Drive** helps minimize mount times by mounting the tape in a physically close drive. For STK drives, a unit of distance is from one Silo/LSM (Library Storage Module) to the next. This means that the tape must go through a pass-through port or elevator. This process is more time consuming. Therefore the default value for STK robots is a higher value which forces the tape to stay in the same Silo/LSM even if it means the drive selected is attached to a controller which is heavily used. All other things being equal, the tape will be mounted in the closest drive.

Shelf Tape Check-In Retry

The number of seconds the PVR will wait before asking the robot if a requested shelf tape has been checked in. The PVR will continue checking at this interval until the tape is checked in. This field applies only if the **Support Shelf Tape** check box is selected. The retry value must be "30" or greater.

Shelf Tape Check-In Alarm

The PVR will periodically log alarm messages when a requested shelf tape has not been checked in. This field specifies the number of minutes between alarms. This field is only active if the **Support Shelf Tape** check box is selected. The alarm value must be "2" or greater.

Dismount Delay

When **Defer Dismounts** is checked, this value is used by the PVL to determine the number of minutes that dismounts are delayed after the last data access.

Retry Mount Time Limit

The default value for this field is "-1". When the default value is used, if an error is encountered during a PVR mount operation, the mount will pend and be retried every five minutes. Setting a value in this field will change the mount behavior to periodically retry the mount until the specified time limit is exceeded. Once exceeded, an error is generated and the mount request is canceled. If the mount request would have resulted in a write operation, the error returned will cause the Core Server to set the **VV Condition** of the associated tape volume to DOWN. Once in DOWN state, the volume will no longer be available for read or write operations. For further

information about the Core Server VV Condition, see *Section 8.5.4.2, “Core Server Tape Volume information window”*.

Drive Error Limit

This field is used in conjunction with the PVR Server **Retry Mount Time Limit**. If the number of consecutive mount errors which occur to any drive in this PVR equal or exceed this value, the drive is automatically locked by the PVL. The only mount errors that apply are those set through the **Retry Mount Time Limit** mechanism. The **Drive Error Count** field in the PVL Drive Information records the number of consecutive errors on a drive by drive basis. To turn off the automatic drive disable feature, set the **Drive Error Limit** to "0" or "-1". Changing its value will not change drive disable behavior until the PVL is recycled.

Defer Dismount Exempt Count

If set greater than "0", this value will specify the number of drives within a PVR that will be exempt from the defer dismount logic. This behavior will take place for all drives within the PVR with a drive/device "Library Unit ID" set (see Drive/Device Configuration). If multiple Library Unit IDs are set, the same Exempt count will apply to all Library Unit ID subsets.

The goal of this option is to allow a finer control of the defer dismount logic to account for transfer costs (for example, elevator, pass through costs). Defer dismount logic to apply logic on a sub library case (unit). The definition of a "unit" is placed in the hands of the administrator and likely based on library characteristics and topology.

Characteristics

Flags for the PVR.

Defer Dismounts

If ON, the PVL will delay the dismounting of a tape cartridge until the drive is required by another job or until the **Dismount Delay** time limit is exceeded. If the **Defer Dismount Exempt Count** is also set, the defer dismount logic will not be exercised on the number of drives set. This may result in free drives available for new mount requests without first generating defer dismount operations. See **Defer Dismount Exempt Count** for details.

A PVL restart is required in order for changes to this flag to take effect.

Support Shelf Tape

If ON, the PVR and the PVL will support the removal of cartridges from the tape library using the **shelf_tape** utility.

ACSLs Packet Version

The packet version used by STK's ACSLS software.

Alternate SSI Port

A site running two SSI (Server System Interface) services on the same platform will specify the non-default port number of the second service via this field. Refer to *the section called “STK PVR additional information”*.

5.3.2.11. RAIT Engine Information window

This window allows you to view and update the type-specific information associated with a RAIT Engine. Any changes made to fields on this window are sent directly to the RAIT Engine after the appropriate button is clicked and are effective immediately.

Field descriptions.

Server Name

The descriptive name of the RAIT Engine.

Number of Request Tasks

The number of RAIT Engine request-processing tasks that currently exist.

Number of Active Requests

The number of active requests that are being handled by the RAIT Engine.

Time of Last Statistics Reset

The time and date when the RAIT Engine statistics were last reset.

Number of Requests Processed

The number of requests that the RAIT Engine has processed since the last time this field was cleared. Click on the **Reset** button to the right of the field to reset the value to zero.

Number of Request Errors

The number of requests that have returned errors since the last time this field was cleared. Click on the **Reset** button to the right of the field to reset the value to zero.

Number of Data Transfers

The number of data transfers that the RAIT Engine has handled since the last time this field was cleared. Click on the **Reset** button to the right of the field to reset the value to zero.

Number of Bytes Moved

The number of bytes of data that the RAIT Engine has transferred since the last time this field was cleared. Click on the **Reset** button to the right of the field to reset the value to zero. (*Note: this field is not implemented in HPSS Release 7.4.1.*)

Total Move Time

The total time, in days, hours, minutes, and seconds, that the RAIT Engine tasks have spent transferring data since the last time this field was cleared. Click on the **Reset** button to the right of the field to reset the value to zero.

Operational notes

1. The RAIT Engine statistics (**Number of Requests Processed**, **Number of Request Errors**, **Number of Data Transfers**, **Number of Bytes Moved**, and **Total Move Time**) may all be reset to zero. If any of these fields are reset the **Time of Last Statistics Reset** field will also be updated with the current time. The statistics fields are all initialized to zero when the RAIT Engine is started.
2. The **Total Move Time** statistic should not be used to determine RAIT Engine throughput rates if more than one request is handled by the RAIT Engine concurrently. Since this is a total for all RAIT Engine tasks involved in the transfer. This total will most likely be considerably greater than the elapsed time necessary to transfer the data.

Related information.

HPSS Error Manual : Problem diagnosis and resolution chapter : Mover problems section.

5.4. Real-Time Monitoring (RTM)

Real-time Monitoring (RTM) is a facility built into the Gatekeeper, the Core Server, the RAIT Engine, and the Mover. Whenever an HPSS thread acting on behalf of a user job or a high-level system activity needs to block or process an operation that may take some time to complete, it will create an RTM record with information about what it is currently doing or waiting for. The RTM utility (**rtmu**) can be used to query the list of the current requests and obtain pertinent data for each request. The RTM requests can be used to construct an overview of the current operations being processed by the HPSS system as well as to monitor the overall health and status of the system. RTM requests that remain in the system for an extended period of time may indicate that the system is under heavy load. If a request remains for an unusually long time, it may indicate a hang condition or a deadlock situation that may require special intervention.

HPSS administrators and operators may use SSM to view the active RTM requests. The *RTM Summary* window lists a summary of the current RTM requests. The *RTM Detail* window displays detailed information for selected RTM requests. Requests in the *RTM Summary* window can be aborted using the **Abort I/O** button. Refer to *Section 5.5, "I/O Abort"*.

5.4.1. RTM Summary List

The screenshot shows a window titled "RTM Summary List" with a menu bar (File, Edit, Column View) and a Help button. It displays a table of RTM request summaries. The table has columns for ID, SubsystemId, Action, State, Time Created, Time Updated, Create Server, Update Server, Rec Cnt, and User Name. The table is sorted by Time Created. The first row is highlighted in blue. The table shows several rows of data, including requests for "read label", "write", "migrate", and "mover read".

ID	SubsystemId	Action	State	Time Created	Time Updated	Create Server	Update Server	Rec Cnt	User Name
0		read label	blocked	00-00:00:12	00-00:00:07	Mover (hpss-...	Mover (hpss-dev-2gr06)	10	
3836530610352750605		write	in progress	00-00:00:14	00-00:00:14	Core Server	Mover (hpss-dev-2gr06)	2	user1
6387102294887170063		write	in progress	00-00:00:21	00-00:00:21	Core Server	Mover (hpss-dev-2gr06)	2	griffith
8784650427408842810		migrate	in progress	00-00:01:56	00-00:00:45	Core Server	Core Server	3	hpssmnp
8784650427408842812		migrate	in progress	00-00:01:56	00-00:00:45	Core Server	Core Server	3	hpssmnp
8784650427408842814		migrate	in progress	00-00:01:56	00-00:00:45	Core Server	Core Server	3	hpssmnp
8784650427408842816		migrate	in progress	00-00:01:56	00-00:00:45	Core Server	Core Server	3	hpssmnp
8784650427408842818		migrate	in progress	00-00:01:56	00-00:00:45	Core Server	Core Server	3	hpssmnp
8784650427408842808		migrate	in progress	00-00:02:18	00-00:00:18	Core Server	Mover (hpss-dev-2gr06)	6	hpssmnp

9/9

Freeze Refresh Dismiss

Detail
Abort I/O
Preferences
Edit
Default

Field descriptions.

RTM Summary List

This is the main portion of the window which displays various information about each RTM request summary.

ID

The RTM request identifier.

Action

The action or operation that this request is currently executing. Examples include "Mover write", "PVL verify", or "write tm" (that is, write a tape mark).

State

The state of the requested operation. Examples include `in progress`, `suspended`, or `blocked`.

SubsystemId

The ID of the HPSS subsystem owning this request.

Time Created

Of all the RTM request records with this ID, this field holds the age of the oldest. The time delta is the age, or time difference between when the request entered the system and the *RTM Summary List* window was last updated.

Time Updated

Of all the RTM requests records with this ID, this field holds the age of the request most recently updated by its server. The time delta is the age, or time difference between when the request was updated by the server and the *RTM Summary List* window was last updated.

Create Server

The server holding the RTM request with the oldest age or Time Created.

Update Server

The server holding the RTM request with the most recent update age or Time Updated.

Rec Cnt

The number of RTM records with this ID.

User Name

The printable username associated with this request (if available).

Path Name

The HPSS pathname of the file associated with this request (if available). The reporting level of the path name is controlled by the environment setting `HPSS_RTM_USE_FULL_PATH`. Refer to the entry for `HPSS_RTM_USE_FULL_PATH` in the *HPSS Installation Guide* Appendix `hpss_env_defs.h` for details.

Wait Reason

A text string indicating what the request is currently waiting for. It may also contain additional information on the resources being waited for.

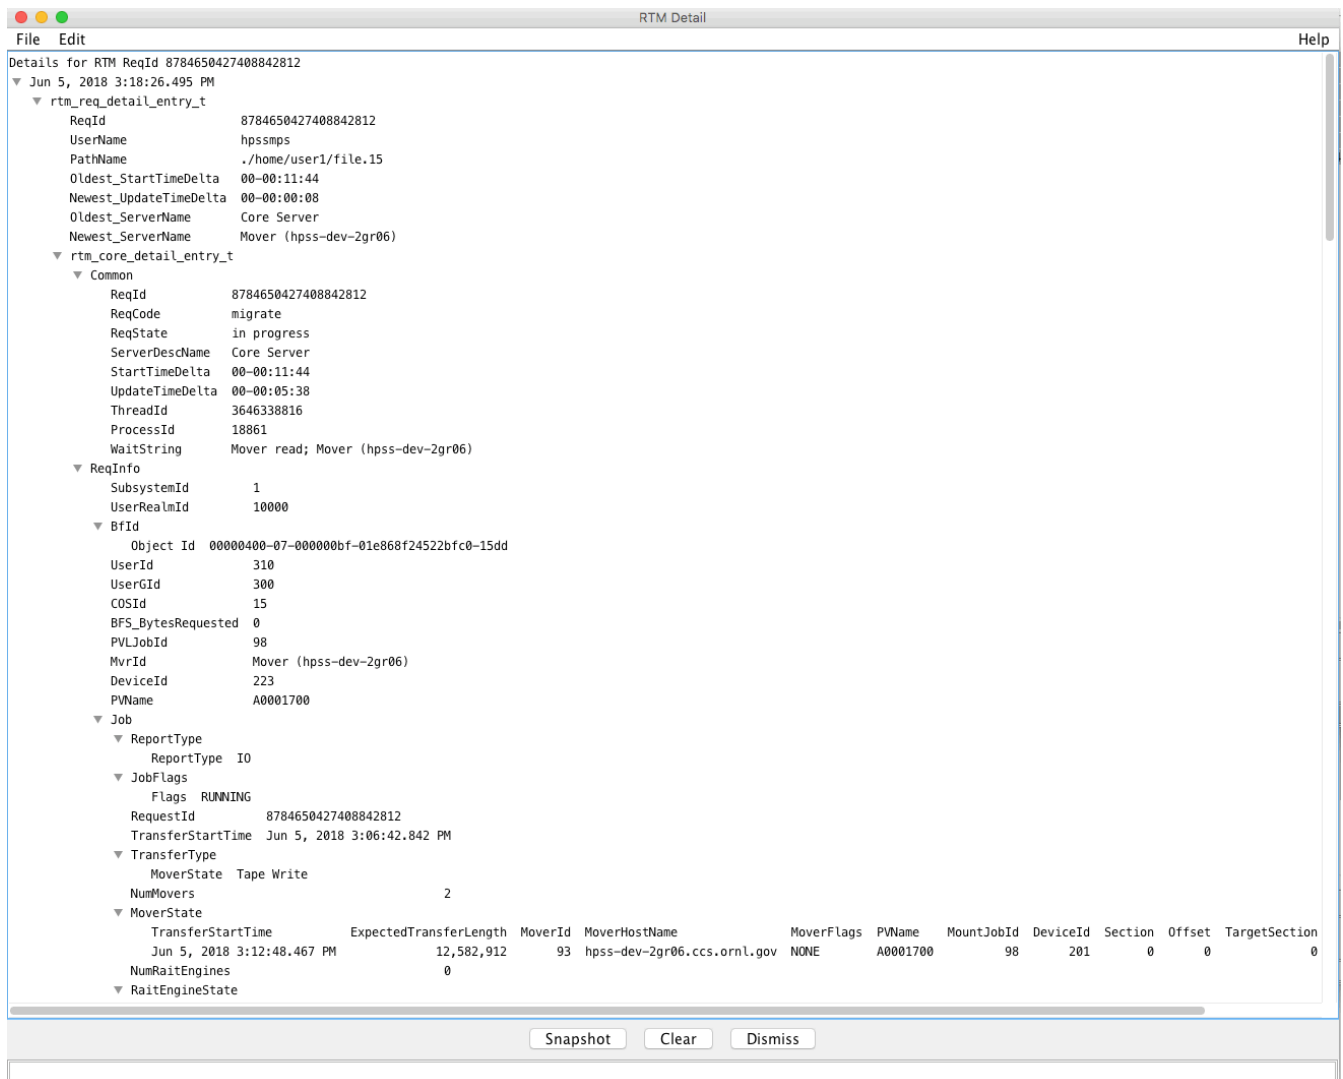
PVLJobId

The PVL job ID (if available).

PVName

The name of the Physical Volume associated with this request (if available).

5.4.2. RTM Detail



The *RTM Detail* window displays a snapshot of the details of the selected RTM requests from the *RTM Summary List* window. This may contain information from multiple servers, Gatekeeper, Core, RAIT Engine, and Mover. The actual data displayed will be different for each server type and is

displayed in a tree structure. Each node of the tree can be expanded or collapsed by clicking the mouse on the tree node indicator. A new snapshot can be taken and added to the display by clicking the **Snapshot** button. The tree can be cleared by clicking the **Clear** button.

Each detailed snapshot taken for an RTM request will contain the date and time the snapshot was taken at the top of the tree. Below the date and time there will be an **rtm_req_detail_entry_t**.

The fields contained in the **rtm_req_detail_entry_t** structure are:

ReqId

The RTM request identifier.

UserName

The printable username associated with this request (if available).

PathName

The HPSS pathname of the file associated with this request (if available). The reporting level of the path name is controlled by the environment setting HPSS_RTM_USE_FULL_PATH. Refer to the entry for HPSS_RTM_USE_FULL_PATH in the *HPSS Installation Guide Appendix hpss_env_defs.h* for details.

Oldest_StartTimeDelta

Of all the RTM request records with this **ReqId**, this field holds the age of the oldest when the snapshot was taken.

Newest_UpdateTimeDelta

Of all the RTM requests records with this **ReqId**, this field holds the age of the request most recently updated by its server when the snapshot was taken.

Oldest_ServerName

The server holding the RTM request with the oldest age or **Oldest_StartTimeDelta**.

Newest_ServerName

The server holding the RTM request with the most recent update age or **Newest_UpdateTimeDelta**.

Each server type that generates RTM request records (Gatekeeper, Core, RAIT Engine, and Mover) has a server-specific RTM data structure which will be displayed as a tree node below the **rtm_req_detail_entry_t** node. These structures are only displayed if there is a server of that type involved in the RTM request. There will be multiple entries if there are more than one server involved in the RTM request.

The server-specific RTM records are called:

- **rtm_gk_detail_entry_t** (Gatekeeper)
- **rtm_core_detail_entry_t** (Core)
- **rtm_mvr_detail_entry_t** (Mover)
- **rtm_rait_detail_entry_t** (RAIT Engine)

The server-specific data structures each have a **Common** and **ReqInfo** structure embedded in them. The **Common** structure will contain the same fields for each server type.

The **Common** fields are:

ReqId

The RTM request identifier.

ReqCode

The action or operation that this request is currently executing. Examples include `Mover write`, `PVL verify`, or `write tm` (that is, write a tape mark).

ReqState

The state of the requested operation. Examples include `in progress`, `suspended`, or `blocked`.

ServerDescName

The descriptive name of the server that holds this RTM request record.

StartTimeDelta

The age of this request since it entered the server. That is, the time difference between when the request entered and when this snapshot was taken.

UpdateTimeDelta

The time difference between when this server last updated this RTM record and when this snapshot was taken.

ThreadId

The ID of the thread processing this request, useful in attaching to the thread to examine its state in the case of a hung request.

ProcessId

The ID of the UNIX process that is processing this request, useful in attaching to the process to examine its state in the case of a hung request.

WaitString

A text string indicating what the request is currently waiting for inside this server.

The **ReqInfo** portions of the server-specific data structures will contain different fields depending on the server type. Not all fields for a particular server-specific **ReqInfo** structure will be displayed. Only those fields that have valid data in them for the particular state of the RTM request will be displayed. The possible fields for each server-specific **ReqInfo** are listed below.

The Core Server **ReqInfo** structure may contain any or none of the following fields:

SubsystemId

The ID of the HPSS subsystem owning this request.

UserRealmId

The realm ID of the user associated with this request.

FilesetId

The ID of the fileset of the file associated with this request.

FilesetName

The text string name of the fileset of the file associated with this request.

BfId

The bitfile ID of the file associated with this request.

BfPathname

The HPSS pathname of the file associated with this request.

UserId

The ID of the user associated with this request.

UserGId

The group ID of the user associated with this request.

COSId

The Class Of Service ID of the file associated with this request.

BFS_SClassId

The storage class ID known to the bitfile server of the file segment associated with this request.

BFS_BytesRequested

The number of bytes to be moved in this request as known by the bitfile server.

BFS_BytesMoved

The number of bytes already moved for this request as known by the bitfile server.

PVLJobId

The ID of the PVL job associated with this request.

MvrId

The ID of the Mover this request is currently waiting on.

DeviceId

The device this request is currently waiting on to complete a data move operation.

Segment

The ID of the storage segment being operated on by this request.

VV

The ID of the virtual volume associated with this request.

PVName

The name of the physical volume associated with this request.

CurrentRelPosition

The current media position of this physical volume as a relative address.

CurrentAbsPosition

The current media position of this physical volume as an absolute address.

ExcessVVsInUse

The tape virtual volumes in use by this storage class have exceeded its quota. Waiting for a this number of virtual volumes to become available.

FamilyId

The ID of the file family of the file associated with this request.

SS_SClassId

The storage class ID known to the storage server of the file/segment associated with this request.

SS_BytesRequested

The number of bytes to be moved by this request as known by the storage server.

SS_BytesMoved

The number of bytes that have been moved for this request as known by the storage server.

Job

A structure which describes the current job tracking information for this request. Its fields are described below:

ReportType

The type of report being delivered. The values are: `IO`, `VV Mount`, `VV Handoff`, `Tape Positioning`, `Tape Alignment`, and `Unknown`.

JobFlags

Flags which describe the status of the job. The values are `RUNNING`, `ABORTING`, and `ABORTED`. The job is in the `ABORTING` state when an abort action is being processed, and once that action has been completed it changes to be `ABORTED`.

RequestId

The request ID.

TransferStartTime

The time that the job began.

TransferType

The type of job being executed.

- `None` (No type)
- `Disk Check` (Check the disk size)
- `Disk Read` (Reading from a disk)
- `Disk Write` (Writing to a disk)
- `Disk Copy Read` (Reading from a disk for a copy)
- `Disk Copy Write` (Writing to a disk for a copy)
- `Disk Mount` (Mounting a disk volume)
- `Tape Read` (Reading from a tape)
- `Tape Write` (Writing to a tape)
- `Tape Verify` (Verify a tape volume)

- Tape Mount (Mounting a tape volume)
- Tape Assignment (Waiting to be assigned a tape)
- Tape Position (Positioning a tape volume)
- Tape Alignment (Aligning tape volume)
- Tape Recommended Access Order (Waiting on an RAO request to a drive)
- Tape Session Queue (Queued session against a volume)
- BFS Scheduling (Waiting on bitfile to be accessible)
- BFS Background Stage Queue (Waiting in the background stage queue)
- BFS Running (Running)

MoverState

Information about a request to a Mover participating in the job. There may be multiple **MoverState** structures included.

TransferStartTime: When the request to the Mover began.

ExpectedTransferLength: The length of the I/O, in bytes.

MoverId: The server ID of the Mover.

MoverHostName: The hostname of the Mover TCP process.

MoverFlags: Information about the state of the I/O request in this Mover.

PVName: The physical volume associated with the Mover.

MountJobId: The PVL job ID.

DeviceId: The device ID being used for the operation.

PVPosition: For positioning requests, the location to position to.

TargetSection: For positioning requests, the section to position to.

AlignmentError: Whether a tape alignment error occurred.

RAITEngineState

Information about a request to a RAIT Engine participating in the job. There may be multiple **RAITEngineState** structures included.

TransferStartTime: When the request to the RAIT Engine began.

ExpectedTransferLength: The length of the I/O, in bytes.

REId: The server ID of the RAIT Engine.

REHostName: The hostname of the RAIT Engine process.

REDescName: The descriptive name of the RAIT Engine.

REFlags: Information about the state of the I/O request in this RAIT Engine.

PVLState

Information about a request to the PVL in the job. There may be multiple **PVLState** structures included.

PVName: The physical volume involved in the operation.

PVFlags: The status of the mount operation (NOT-MOUNTED, MOUNT, MOUNTED, DEALLOC, ACTIVE).

PVLJobId: The PVL job ID.

BFSState

Information about the BFS state for the job.

BfId: The bitfile identifier involved in the job.

JobDescription: A string describing what the job is doing in the BFS.

PendingJobId

PVL job ID.

The Mover **ReqInfo** structure may contain any or none of the following fields:

InitialTransferOffset

The offset of the start of the Mover device for the transfer associated with this request.

CurrentTransferOffset

The current offset of the Mover device for the transfer associated with this request.

DeviceId

The ID of the device associated with this request.

VolumeId

The name of the volume associated with this request.

BytesRequested

The number of bytes the Mover has been asked to transfer for this request.

BytesMoved

The number of bytes the Mover has already transferred for this request.

The RAIT Engine **ReqInfo** structure may contain any or none of the following fields:

InitialTransferOffset

The offset of the start of the transfer associated with this request.

CurrentTransferOffset

The current offset of the transfer associated with this request.

BytesRequested

The number of bytes the RAIT Engine has been asked to transfer for this request.

BytesMoved

The number of bytes the RAIT Engine has already transferred for this request.

ParityBufferIndex

The index number of the buffer that the parity thread is currently computing.

The Gatekeeper **ReqInfo** structure may contain any or none of the following fields:

AuthorizedCaller

Indicates whether the caller is authorized or not.

BitFileId

The bitfile ID of the file associated with this request.

ConnectionId

The ID of the (Core Server) connection associated with this request.

ClientConnectedId

The ID of the end client connection associated with this request.

ControlNo

A unique control ID used by the Gatekeeper server.

RealmId

The realm ID of the user associated with this request.

GroupId

The group ID of the user associated with this request.

HostAddr

The address of the originating host.

RequestType

The type of this request (Open, Create, or Stage).

Oflag

The Open flags associated with this file open request.

StageFlags

Flags associated with this file stage operation.

StageLength

The number of bytes to stage.

StageOffset

The offset of the file where the stage is to begin.

StageStorageLevel

The storage class level that the file will be staged to.

UserId

The ID of the user associated with this request.

WaitTime

The length of time this request has been waiting.

5.5. I/O Abort

This section describes how to abort I/O operations, and what to expect from HPSS following an abort request.

5.5.1. Aborting an I/O

I/O Abort applies to a number of HPSS operations; however, there are many operations which may be beyond its scope. The following types of requests can be aborted:

- Application I/O requests to disk or tape via any client interface
- Stages, including background stages and queued background stages
- Open requests which require a stage
- Migration operations including migrations from MPS or Force Migration requests from Recover
- Repack operations
- Logical Block Protection (LBP) Verify operations
- Requests for tape access waiting in the HPSS tape scheduler

Here are some examples of requests which fall outside of this capability and which cannot be aborted via this feature:

- A change COS request; the I/O can be canceled but the change COS request remains
- A long running name space operation such as a directory listing or a UDA query
- A PVR cartridge move operation
- A request to get file attributes which has become blocked due to a lock on the file
- Disk purge operations
- Creating or deleting HPSS storage resources

Requests for I/O can be aborted in several ways:

- Via the SSM GUI *RTM Summary* window
- Via the ADM command line interface
- Via the SSM GUI *Active Storage Classes* window (for migrations)
- Via applications that support aborting I/O via the Client API

To abort a request from the GUI, simply select one or more requests from the *RTM Summary* window, select the "Abort I/O" button, and confirm.

To abort a request from the ADM, use the **rtm cancel** command:

```
hpssadm> rtm cancel -id <request id>
```

Request IDs can be identified using the RTM summary or detail information with the HPSS ADM, GUI, or `rtmu`.

An abort request will return successfully when the HPSS Core Server has received the abort request and taken some action to cancel the I/O request. The abort request will return prior to the request being completed and leaving the RTM window.

In some cases, a request may be in a state where it is not abortable and will return with an error. This is likely due to the request being very early in its life and not registered as an abortable job, or not being a request which can be aborted because it is not related to an I/O request (for example, a Core Server tape handoff operation).

Issuing a **Suspend** or **Stop** on an active storage class migration will abort ongoing I/O. This will allow HPSS administrators to quickly suspend/stop migrations for maintenance operations.

5.5.2. Expectations in Aborting I/O

The primary use case for I/O Abort is in dealing with stuck or hung jobs; jobs which may never complete. Another significant use case is ending accidental work that has been spun up against the system; for example, an accidental large stage request, which may occupy the system for a significant amount of time for no purpose. For these use cases, I/O Abort tends to work very well. Because small I/O requests can very quickly enter and exit the system, it can be difficult to cancel ongoing small file I/O; however, this is not a situation that I/O Abort is intended to address. There are other methods for dealing with this.

Another important use case for I/O Abort is quickly quiescing migration activity in preparation for a downtime, or other maintenance work, by issuing a **Suspend** or **Stop** on each active storage class migration.

HPSS implements its abort interface as a "lazy" interface. The request to abort returns once the Core Server has received it and determined whether the operation associated with the incoming request ID can be aborted. This lazy method is employed because, depending upon the current state of the request, it could take several minutes for the request to ultimately wind down. For example, if a request for a stage is aborted and the tape is in the middle of being mounted the tape mount may still have to complete and then be dismounted before the request can end, which can take a significant amount of time.



It is important to consider that aborting a request may generate significant work, such as dismounting a large number of tape volumes which were in use by the aborted request. This could temporarily delay other work attempting to use the same library for other purposes.

Since the abort interface is "lazy", the request will continue to appear in RTM. However, in the case that some action has been taken to abort the request, the RTM details will list the job as being in an "ABORTING" or "ABORTED" state, depending on whether the abort action has already been taken or not.



Requests will remain in the *RTM Summary* window and continue to be updated until they are successfully aborted or they complete.

Aborted requests are logged differently in many cases. The abort request is logged as an EVENT message for every request that is aborted. In addition, resulting log messages that would normally be ALARM messages will generally be reduced to DEBUG messages.

From an HPSS client application perspective, an HPSS I/O request that has been aborted will appear as an EIO (Input/Output Error). Unlike some errors, the Client API will not retry an aborted request automatically. Client applications can retrieve more detailed information about the result from via the Client API and thereby distinguish a failed request from an aborted one.

5.6. Starting HPSS

This section describes how to start and stop HPSS servers and prerequisite software.

5.6.1. Starting HPSS prerequisite software

Invoke the `/opt/hpss/bin/rc.hpss` script as root to start the HPSS prerequisite software such as DB2, Kerberos and LDAP. By default, stdout and stderr from commands started by **rc.hpss** are redirected to the logfile `HPSS_PATH_LOG/hpss_stdout_stderr.log` unless the **-o** option is used. Specifying the **-o** option to **rc.hpss** will cause stdout and stderr of commands invoked within the script to be sent to `/dev/console`, the terminal window where **rc.hpss** is running.

To start the prerequisite software, use the **-p** option to **rc.hpss**:

```
% su -  
% /opt/hpss/bin/rc.hpss -p [start]
```

5.6.2. Starting HPSS servers

Most HPSS servers are started via SSM which calls an appropriate Startup Daemon to begin execution of the selected HPSS server. This means that SSM and the Startup Daemons must be started in a different way. This section covers how to start the Startup Daemons, how to start SSM, and how to start the other HPSS servers once SSM and the Startup Daemons are up and running.

5.6.2.1. Starting the Startup Daemons

Invoke the `/opt/hpss/bin/rc.hpss` script as root to start the HPSS Startup Daemon. By default, both stdout and stderr are redirected to the logfile `HPSS_PATH_LOG/hpss_stdout_stderr.log` and not to `/dev/console` since this can have a negative impact on performance. Specifying the **-o** option will cause redirection of output to `/dev/console`.

To start the Startup Daemon, use the **-d** option to **rc.hpss**:

```
% su -  
% /opt/hpss/bin/rc.hpss -d [start]
```

5.6.2.2. Starting SSM

The SSM System Manager configuration metadata should have already been created by **mkhpss** as part of the infrastructure configuration. After SSM is started, this configuration metadata may be modified if necessary. Refer to the *HPSS infrastructure* section of the *HPSS Installation Guide* (in Chapter *HPSS basics*) for more information.

The SSM server startup script, **rc.hpss**, can be used to invoke the HPSS System Manager. The server requires a number of environment variable values to be used as its command line arguments and to be used to configure and manage the HPSS servers. These variables are defined by the system. The system defined values can be overridden in the `env.conf` file. They should have been edited to reflect the site's configuration during the HPSS infrastructure configuration phase.

Before starting up the System Manager, ensure that DB2 and other prerequisite software are up and running.

The **rc.hpss** script will not start the System Manager if it is already running on the same node. Ensure that only one instance of the SSM System Manager is running at any one time.

To start the SSM System Manager, use the **-m** option to **rc.hpss**:

```
% su -  
% /opt/hpss/bin/rc.hpss -m [start]
```

If the SSM's server name has changed since `mkhpss` was run, the environment variable `HPSS_DESC_SSMSM` must be set to match the new server name.

For best performance, it is recommended that the SSM graphical user interface (GUI) be installed on the SSM user's desktop machine. However, authorized SSM users can invoke the **hpssgui** script directly from the HPSS server machine if desired. However, most users of the SSM command line interface (**hpssadm**) will most likely invoke the script on the server machine.

5.6.2.3. Starting other HPSS servers

Other servers are started using SSM or can be automatically started upon startup of the System Manager by using the Core Server **Auto Startup** feature. This feature can be set using the *Core Server Configuration* window. Refer to *Section 5.2.1.2, "Execution controls"*.

When starting a server from SSM using the *Servers* window, one or more HPSS servers can be started by the same request. SSM forwards the server's configuration data to the appropriate HPSS Startup Daemon. The Startup Daemon invokes the server with the appropriate input arguments. After the server is started, SSM attempts to establish a connection with the server. The server status is reported in the **Status** field on the *Servers* window. Refer to *Section 5.1, "Server list"* for more information.

Before starting a server, ensure that DB2 and other prerequisite software are running, Kerberos is running (if using Kerberos authentication), that the Startup Daemon on the server's node is running and has a connection with SSM, and that the server is configured and its **Executable** flag is set. Ensure that any required PVR controllers are up and running before bringing up the PVRs.

To start a server, select one or more desired servers from the *Servers* window and click on the **Start** button. Verify the result of the request in the message area on the *Servers* window. In addition, monitor the *Alarms and Events* window for the "Server Initialized" event.

The Startup Daemon allows only one instance of a configured server to be brought up at a time. If a server is already running, the subsequent startup request for that server will fail.

5.6.2.4. Automatic server restart

HPSS servers may be configured to automatically restart. An **Auto Restart Count** is contained in each server configuration. This count may be used to define the number of times a failed server will

be automatically restarted without any administrative action. The **Auto Restart Count** is ignored for the Startup Daemon and the SSM System Manager.

A failure is defined as a server terminating abnormally, and an abnormal termination is the result of the server exiting with a nonzero return code or as a result of a signal other than SIGTERM or SIGKILL. Termination due to an administrative halt or shutdown is not considered to be an abnormal termination and a server will not be automatically restarted.

The following **Auto Restart Count** values may be set:

Table 5.2. Auto Restart Count values

Auto Restart Count value	Meaning
0	Do not restart
-1	Infinite restart attempts
n	Restart a failed server n times (where $n > 0$)

After the **Auto Restart Count** is exhausted, the Startup Daemon will not attempt to restart the server. If the server executes for more than one hour without failing, the number of failures will be set back to zero.



*Do not confuse this feature with the Core Server **Auto Startup** feature. This feature can be set using the Core Server Configuration window described in Section 5.2.1.2, “Execution controls”.*

5.7. Stopping HPSS

5.7.1. Shutting down an HPSS server

One or more HPSS servers can be shut down in a single SSM request. SSM will send each server a request to perform a controlled shutdown. The server will attempt to complete its current tasks, discontinue accepting new tasks, and inform SSM of its termination prior to exiting.

To initiate the request, select the desired servers on the *Servers* window and click on the **Shut Down** button. SSM will confirm the request through a pop-up window. The SSM user may wish to verify the result of the request in the status bar on the *Servers* window as well as to monitor the *Alarms and Events* window for the "Server Terminated" event.



*HPSS Startup Daemons and the SSM System Manager cannot be shut down from the *Servers* window. Select **System Manager** from the **Shutdown** submenu of the **Operations** menu of the HPSS Health and Status window to shut down the System Manager. Use the **rc.hpss** script stop option to shut down either the System Manager or the Startup Daemon.*

Servers may not terminate immediately since they may wait for pending operations to complete before terminating. Monitor the *Alarms and Events* window for possible notifications of delay from the server. In addition, monitor the *Servers* window for the server's shutdown status. If an immediate

shutdown is required, use the **Force Halt** operation described in *Section 5.7.3, “Halting an HPSS server”*.

SSM must have a connection to the Startup Daemon associated with the server or the shutdown operation will fail.

5.7.2. Shutting down all HPSS servers

Most HPSS servers (all but SSM and the Startup Daemon) can be shut down by selecting **Shut Down All Non-SSM Servers** from the **Shutdown** submenu of the **Operations** menu. SSM will process this request similarly to the **Shutdown Server** request issued from the *Servers* window.

This feature is useful when the entire HPSS system must be shut down, leaving SSM up to monitor the server shutdown process. When the servers are down, the SSM server and sessions can be shut down by selecting **System Manager** from the **Shutdown** submenu of the **Operations** menu. The SSM client sessions will detect that the SSM System Manager has been shut down and will then inform the user of the event.

5.7.3. Halting an HPSS server

Halting an HPSS server stops it without giving it an opportunity to shut down gracefully. When the command is received by the server, the server simply terminates immediately without doing any of its normal shutdown processing. It should be used only if it is necessary to shut down the server as quickly as possible or if attempts to shut the server down using the **Shutdown** button have failed.

One or more HPSS servers can be halted in a single request. SSM will issue a request to each selected server's Startup Daemon to issue a SIGKILL signal to its HPSS server.

To initiate the request, select the desired servers on the *Servers* window and click on the **Force Halt** button. SSM will confirm the request through a pop-up window. Verify the result of the request in the status bar on the *Servers* window. In addition, monitor the *Alarms and Events* window for the "Server Halted" alarm.

HPSS Startup Daemons and the SSM System Manager cannot be halted from the *Servers* window.

SSM must have a connection to the Startup Daemon associated with the server or the Halt Server operation may fail.

5.7.4. Shutting down the SSM server

To shut down the SSM server, select **System Manager** from the **Shutdown** submenu of the **Operation** menu on the *HPSS Health and Status* window. All SSM user sessions, including graphical and command line sessions, will detect that the System Manager has exited. Choosing this option will pop up a confirmation window which allows the shutdown request to be approved or canceled.

As the System Manager exits, a notification window will pop up on each logged-on SSM graphical user interface session informing the user that the GUI has lost connection to the System Manager. The session user can click the **Exit** button on this window to exit the SSM session or the **Login** button to log out of the session and get a new login window, which can be used once the System Manager has been restarted.

At each command line user interface session, a notification message will be displayed that the client connection to the System Manager has been lost and the session will be terminated.

These notifications will also happen whenever the client is having communication problems with the System Manager, even if the System Manager is still running.

Additionally the System Manager can be shut down with the **-m** option of **rc.hpss**:

```
% su -  
% /opt/hpss/bin/rc.hpss -m stop
```

5.7.5. Shutting down the Startup Daemon

To stop the Startup Daemon, use the **-d** option of **rc.hpss**:

```
% su -  
% /opt/hpss/bin/rc.hpss -d stop
```

5.7.6. Stopping the prerequisite software

To stop the prerequisite software:, use the **-p** option of **rc.hpss**:

```
% su -  
% /opt/hpss/bin/rc.hpss -p stop
```

This will stop Kerberos, LDAP and DB2 if the corresponding shutdown scripts exist in `/var/hpss/` etc.

5.8. Server repair and reinitialization

This section describes how to repair or reinitialize an HPSS server.

5.8.1. Repairing an HPSS server

Repairing a server is the process of instructing the server to reset its internal state and status variables to their nominal settings. These variables include Operational State, Usage State, Administrative State, Service Status, Security Status, Software Status, Hardware Status, and Communication Status. Operational State is displayed prominently on the *Servers* window. If a server encounters an error as it runs, it may set one or more of these state and status variables to a value indicating the type and severity of the error. Alarm messages may be generated as well. Repairing a server simply resets all of these state and status variables to their nominal settings. One or more servers can be repaired with each request.

To initiate the request, select the desired servers from the *Servers* window and click on the **Repair** button. Verify the result of the request in the status bar on the *Servers* window. In addition, monitor the *Alarms and Events* window for the "Server Repaired" event.

Repairing a server does not correct the underlying problem that caused the server's reported state to change. Rather, it is a means for the administrator to notify the server that the underlying problem has been corrected or dismissed. It is an "alarm reset".

5.8.2. Reinitializing a server

Reinitializing a server causes it to reread all or part of its configuration metadata and reset its internal representation of those values. In many cases reinitialization or clicking the **Apply Config** button is all that is needed to cause a server to start using a new configuration. In other cases it may be necessary to restart a server to cause it to use a new configuration.

HPSS servers vary widely in their support of the reinitialization function. The following table lists the behavior of each HPSS server when it is told to reinitialize.

Table 5.3. Server reinitialization behavior

Server	Behavior
Location Server	This server supports full reinitialization. This means that reinitialization is as effective as restarting the server.
Core Server	<p>Resets the Maximum Active I/O Requests, Maximum Active Copy Requests, COS Copy to Disk, COS Change Retry Count, Tape Dismount Delay, Tape Handoff Delay, PVL Max Connection Wait and Fragment Smallest Block values to the values in the specific configuration metadata record.</p> <p>Note that reducing the Maximum Active I/O Request or Maximum Active Copy Requests does not result in any running operations being quiesced - use I/O Abort to reduce the total number of running operations.</p> <p>Sets the server's diagnostic Logging Options to the values in the specific configuration metadata record.</p> <p>Reloads cached Class of Service information for the system.</p> <p>Reloads cached Migration Policy for those policies that were already cached in the server's memory. Does not add new policies to, or remove deleted policies from the cache.</p> <p>Reloads cached global Root User ID, Root is Superuser, Can change UID to self if has Control Perm, Can change UID if has Delete Perm on Security ACL, Object names can contain unprintable characters, Trashcans Enabled, Number of Trashcan Threads, Trash Incinerator Interval, Trash Unlink Eligible Time, and Trash Statistics Interval values using values read from the global configuration metadata.</p> <p>Reloads cached global and subsystem Metadata Space Warning Threshold, Metadata Space Critical Threshold, Metadata Space Monitor Interval, DB Log Monitor Interval, and COS Change Stream Count</p>

Server	Behavior
	<p>values using values read from the global and subsystem configuration metadata.</p> <p>Reloads cached Default Class of Service and Gatekeeper values using values read from the subsystem configuration metadata.</p> <p>Releases the in-memory cache of File Family IDs, allowing it to be reconstructed. If a file family has been removed from the list of defined families, this step causes it to be unknown to the Core Server.</p> <p>Reinitializes the in-memory cache of environment variable values - rereads the <code>env.conf</code> file.</p> <p>Adjusts the maximum number of open files supported by the Core Server when the number of migration streams increases.</p> <p>Refreshes the Core Server's list of available RAIT Engines.</p> <p>Refreshes the Core Server's Manage Object with the new values from this list.</p> <p>Refreshes the Core Server's cached Mover and RAIT Engine configuration metadata. If a Mover or RAIT Engine's configuration metadata is changed, reinitializing the Core Server will make the Core Server aware of the changes.</p> <p>Closes all unused disk Mover communication sockets.</p> <p>Rereads the accounting policy and refreshes the account validation cache.</p> <p>Update server basic configuration thread, request queue, and connection configuration.</p> <p>Any Core Server configuration information not listed above is not changed when the server is reinitialized. For example, security-related configuration settings.</p>
SSM	This server rereads the environment. Update server basic configuration thread, request queue, and connection configuration.
Mover	Rereads log policy only. Update server basic configuration thread, request queue, and connection configuration.
PVR	Reinitializes the internal drive list cache. For the SCSI PVR, also reinitialize the control path list. Update server

Server	Behavior
	basic configuration thread, request queue, and connection configuration.
MPS	Reloads the Storage Class Update Interval , Core Server API Failures Allowed , and MPS Report File Name type-specific configuration values.
GK	Reloads the <i>Accounting Policy</i> Options flags.
All other servers	Update server basic configuration thread, request queue, and connection configuration.

To reinitialize a server, select the desired servers from the *Servers* window and click on the **Reinitialize** button. SSM will confirm the request through a pop-up window. The request results may be verified via the status bar on the *Servers* window as well through the *Alarms and Events* window.



If a configuration setting has been changed and that change can be applied through reinitialization, the **Apply Config** button on the *HPSS Health and Status* screen will become enabled. Clicking the **Apply Config** button will reinitialize all of the servers, except the Startup Daemon and the System Manager, and it is the recommended method to ensure the setting is applied.

If a server does not support reinitialization, a pop-up window will be displayed to inform the user of the invalid request. Warnings are not displayed for servers that only partially support reinitialization. Servers that do not support reinitialization, or those that do not support reinitializing the settings in question, must be restarted in order for configuration modifications to take effect.

Some groups of servers depend on consistent configuration information to run properly. For example, the Core Server and Migration/Purge Server must agree on Class of Service, hierarchy, and storage class configurations. When changes are made to these sorts of configuration information, all of the servers that rely on these settings should be restarted to make sure they agree. In most cases, as mentioned in the note above, the **Apply Config** button can be used and is the recommended method to ensure these configuration settings are applied across all of the servers.

5.9. Forcing an SSM connection

When the status of the server's connection with SSM is UP/UNCONNECTED, select the server and then click on the **Force Connect** button on the *Servers* window. SSM will attempt to establish the connection with the server and reflect the latest data in the **Status** field on the *Servers* window.

The SSM System Manager periodically retries the connection with all the servers that are configured to execute. The **Force Connect** button may be used to trigger an immediate connection request.

Chapter 6. Storage configuration

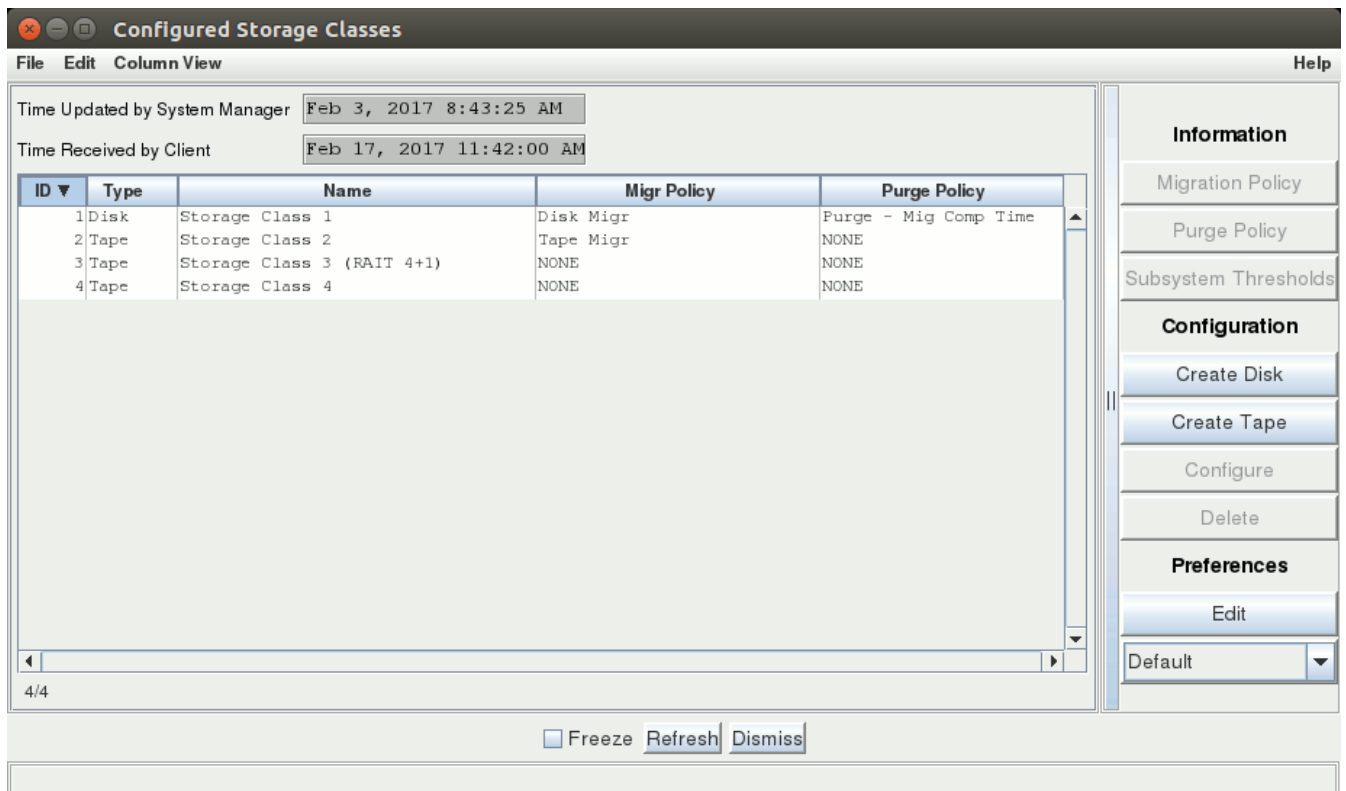
This chapter describes the procedures for creating, modifying, and deleting storage classes, hierarchies, Classes of Service, migration policies, purge policies, and file families.

6.1. Storage classes

This section describes the process of configuring storage classes.

6.1.1. Configured Storage Classes window

A storage class can be created and managed using the *Configured Storage Classes* window.



This window displays all storage classes in the HPSS system, both empty and non-empty. This window is displayed by first selecting **Configure** from the *HPSS Health and Status* window. This will display a drop-down menu. Select **Storage Space** from this menu. This will display another drop-down menu. Select **Storage Classes** from this menu.



When adding or changing a storage class, the Core and Migration/Purge Servers must be reinitialized or restarted in order for the change to take effect. It is recommended to use the **Apply Config** button on the *HPSS Health and Status* window to reinitialize the servers.

See also the related window *Active Storage Classes*, as described in Section 8.3.1, “Active Storage Classes window”. The *Active Storage Classes* window does not list any empty storage classes, but only those to which volumes have been assigned.

Field descriptions.

Storage Classes List

The central panel in this window shows the list of configured storage classes. The columns display selected fields from the storage class configuration windows discussed in the following sections.

Information buttons.

Migration Policy

Opens the configuration window for the migration policy that is configured for the selected storage class. This button will be disabled if no storage classes are selected in the storage classes list or the selected storage class does not have a migration policy.

Purge Policy

Opens the configuration window for the purge policy that is configured for the selected storage class. This button will be disabled if no storage classes are selected in the storage classes list or the selected storage class does not have a purge policy.

Subsystem Thresholds

Opens either a disk or tape subsystem threshold configuration window, as appropriate, for the selected storage class.

Configuration buttons.

Create Disk

Opens a *Disk Storage Class Configuration* window containing default values for a new disk storage class.

Create Tape

Opens a *Tape Storage Class Configuration* window containing default values for a new tape storage class.

Configure

Opens the selected storage class configurations for editing.

Delete

Deletes the one or more selected storage classes.

6.1.2. Disk Storage Class Configuration

Disk Storage Class Configuration

File Edit Help

Storage Class ID: 5

Storage Class Name: Disk - Storage Class 5

Storage Class Type: Disk

Migration Policy: 5 | Migration Policy 5

Purge Policy: 3 | Purge Policy 3

Warning Threshold: 80 percent

Critical Threshold: 90 percent

Optimum Access Size: 0

Average Number of Storage Segments: 4

Storage Segment Size

Media Type: Generic - Default Disk

Media Block Size (MBS): 4KB

VV Block Size (VVBS): 1MB

Data Stripe Width (DSW): 1

Data Stripe Length (DSL): 1MB (VVBS * DSW)

PV Size (PVSIZE): 512MB

VV Size (VVSIZE): 512MB (PVSIZE * DSW)

Min Multiplier (MINMULT): 1

Min Storage Segment Size (MINSEG): 1MB (MINMULT * DSL)

Max Multiplier (MAXMULT): 1 (2^0)

Max Storage Segment Size (MAXSEG): 1MB (MAXMULT * MINSEG)

Min Segments in VV (MINSVV): 512 (VVSIZE / MAXSEG)

Max Segments in VV (MAXSVV): 512 (VVSIZE / MINSEG)

Transfer Rate

Device I/O Rate (DIOR): 3072 KB/sec

Stripe Transfer Rate (STR): 3072 KB/sec (DSW * DIOR)

Update Delete Start Over Clone (partial) Clone (full) Dismiss

This window is used to manage disk storage class configurations.

Field descriptions.

Storage Class ID

The numeric identifier assigned to the storage class.

Storage Class Name

The descriptive name of the storage class.

Storage Class Type

The type of the storage class (always `Disk`).

Migration Policy

The migration policy associated with this storage class, or `None` if no migration is desired.



Advice

Do not configure a migration policy for a storage class at the lowest level in a hierarchy.

*If a migration policy is added to a storage class after files are created in the storage class, those files may never be migrated. Use the **mkmprec** utility to correct this problem. See the **mkmprec** man page for more information.*

Purge Policy

The purge policy associated with this storage class, or `None` if no purge is desired.



Advice

Do not configure a purge policy for a tape storage class or for any storage class which does not have a migration policy in effect. Purge policies only apply to disk storage classes. Purging from tape is managed as a special characteristic of the tape migration policy.

Warning Threshold

A threshold for space used in this storage class expressed as a percentage of the total space defined in the storage class. Alarms will be sent to SSM periodically when free space in the storage class falls below this value, and the **Space Thresholds** field on the *HPSS Health and Status* window will be changed to `Warning`. Note that this field will not have any effect if overridden by storage subsystem-specific thresholds.

Critical Threshold

Another threshold for space used in this storage class expressed as a percentage of the total space defined in the storage class. Alarms will be sent to SSM periodically when free space in the storage class falls below this value, and the **Space Thresholds** field on the *HPSS Health and Status* window will be changed to `Critical`. Note that this field will not have any effect if overridden by storage subsystem-specific thresholds.



Advice

These values determine limits beyond which the MPS sends alarms indicating that storage space is running low. The thresholds on disk are measured in bytes and expressed in terms of the percentage of total space in the storage class. The threshold may be disabled by setting the value to 100 percent.

Optimum Access Size

The optimal transmission size to be used for a transfer request using this storage class. (Not currently used by HPSS. May be used by site-specific applications.)

Average Number of Storage Segments

The **Average Number of Storage Segments** is used by the Core Server as an aid in picking the size of disk storage segments when writing a file to disk. This field is ignored when the storage class is at the top level of the hierarchy in a COS which uses Variable Length or Fixed Length Max Style segment allocation. See the discussion of the **COS Allocation Method** field and **Truncate Final Segment** flag in *Section 6.3.2, "Class of Service Configuration window"* for details on segment allocation methods.

When the storage class is at the top level of a COS which uses Fixed Length Classic Style segment allocation, or when the storage class is at any level below the top, regardless of the COS segment allocation method, the Core Server requires that all of the disk storage segments occupied by a given file be the same size, and must be within the range of sizes given by the Minimum Storage Segment Size and Maximum Storage Segment Size parameters.

If **Max Storage Segment Size (MAXSEG)** is greater than **Min Storage Segment Size (MINSEG)**, the Core Server attempts to select a storage segment size for a file such that the file can be stored in **Average Number of Storage Segments** segments. If the **Max Storage Segment Size (MAXSEG)** is equal to the **Min Storage Segment Size (MINSEG)**, the Core Server will use this value for the storage segment size, and the file will occupy as many of these segments as necessary.

The principal effect of changing the value of the **Average Number of Storage Segments** is on disk volume fragmentation. If the files to be stored in a storage class are large relative to the size of the disk VVs, fragmentation of the volumes may make it difficult to find space for a new segment. Setting **Average Number of Storage Segments** to a larger value will increase the number of segments occupied by files, and decrease the segment size. Fragmentation of the volumes will be reduced, but the amount of metadata required to describe the files will be increased.

Conversely, if files are small relative to the size of the disk VVs, smaller values of **Average Number of Storage Segments** increase the size of the storage segments, and decrease the number of segments occupied by each file. This reduces the metadata storage requirements of the file.

The number of segments in small HPSS files can have a significant impact on transfer performance. To maximize the transfer performance, set this value to "1". Keep in mind that this will result in less effective disk space utilization. On average, you will use only 50% of your disk space with this selection.

Storage Segment Size

Media Type

The media type of all volumes in the storage class.

Media Block Size (MBS)

The **Media Block Size** is the size in bytes of a physical data block on the media. For disk, the value must be a multiple of the physical block size used by the disk hardware. For example, if the disk hardware stores data in 512-byte sectors, 2048 would be a valid entry in this field, but 2000 would not.



Advice

*The **Media Block Size** should be set to a value appropriate for the volume type. See the Media block size selection section of the HPSS Installation Guide (in the HPSS planning chapter) for some recommendations.*

VV Block Size (VVBS)

The virtual volume block size is the number of bytes written to an element of a striped VV before the data stream moves to the next stripe element. It can be thought of as the stride length of striped data transfer operations. The length of the VV block has an effect on the striping efficiency. Short VV blocks cause more protocol overhead when writing striped devices. In non-striped applications, VV Block Size has little meaning so any convenient multiple of the Media Block Size will do.

**Advice**

When choosing a VV Block Size, the administrator should consider the characteristics of any data source or sink that will be copied to or from. Best performance of striped copies usually occurs when the VV Block Sizes of the source and sink are equal. This minimizes the data movement protocol overhead and helps to keep the data streams flowing smoothly.

VV Block Size must be an integer multiple of the Media Block Size.

See the Virtual volume block size selection (disk) and the Virtual volume block size selection (tape) sections of the HPSS Installation Guide (in the HPSS planning chapter) for information.

Data Stripe Width (DSW)

The number of physical volumes in a virtual volume in this storage class.

Data Stripe Length (DSL)

The **Data Stripe Length (DSL)** is the product of the **VV Block Size (VVBS)** and the **Data Stripe Width (DSW)**.

PV Size (PVSIZE)

The size of the disk volume in bytes.

VV Size (VVSIZE)

The virtual volume size. The product of the **PV Size (PVSIZE)** and the **Data Stripe Width (DSW)**.

Min Multiplier (MINMULT)

The **Min Storage Segment Size (MINSEG)** field is set to the product of this value and **Stripe Length (SL)**.

Min Storage Segment Size (MINSEG)

The lower bound for storage segment sizes created on volumes in this storage class. This value is the product of the **Stripe Length (SL)** and the **Min Multiplier (MINMULT)**.

Max Multiplier (MAXMULT)

The **Max Storage Segment Size (MAXSEG)** must be a power of two multiple of the **Stripe Length (SL)**. This selection list contains the valid power of two values from 1 (2^0) to 16,777,216 (2^{24}). Select the appropriate multiplier from the selection list. The field **Max Storage Segment Size (MAXSEG)** is set to the product of this value and **Min Storage Segment Size (MINSEG)**.

Max Storage Segment Size (MAXSEG)

The upper bound for storage segment sizes created on volumes in this storage class. This must be a power of two multiple of the **Stripe Length (SL)**. This value is the product of the **Min Storage Segment Size (MINSEG)** and the **Max Multiplier (MAXMULT)**.

Min Segments in VV (MINSVV)

The minimum number of segments in a virtual volume. This value is calculated as the **VV SIZE (VVSIZE)** parameter divided by the **Max Storage Segment Size (MAXSEG)** parameter.

Max Segments in VV (MAXSVV)

The maximum number of segments in a virtual volume. This value is calculated as the **VV SIZE (VVSIZE)** parameter divided by the **Min Storage Segment Size (MINSEG)** parameter.

Transfer Rate

Device I/O Rate (DIOR)

The approximate data transfer speed, in kilobytes per second (KB/sec), which can be achieved by devices corresponding to **Media Type**. In disk storage classes, this field is informational only.

Stripe Transfer Rate (STR)

The estimated aggregate transfer rate for volumes of this type. This value is the product of the **Data Stripe Width (DSW)** and the **Device I/O Rate (DIOR)**.

6.1.3. Tape Storage Class Configuration

The screenshot shows the 'Tape Storage Class Configuration' window with the following settings:

- Storage Class ID:** 9003
- Storage Class Name:** Storage Class 9003
- Storage Class Type:** Tape
- Migration Policy:** NONE
- Warning Threshold:** 10 volumes
- Critical Threshold:** 5 volumes
- Optimum Access Size:** 0
- Average Latency:** 20 seconds
- Maximum VVs To Write:** 10

Storage Segment Size

- Media Type:** IBM - 3580 (LTO) Gen8 Tape
- ☐ Block Protection
- Media Block Size (MBS):** 512KB
- VV Block Size (VVBS):** 1MB
- Data Stripe Width (DSW):** 1
- Data Stripe Length:** 1MB (VVBS * DSW)
- PV Size (PVSIZ):** 12000GB
- VV Size (VVSIZ):** 12000GB (PVSIZ * DSW)
- Max Storage Segment Size:** 0

Transfer Rate

- Device I/O Rate (DIOR):** 307200 KB/sec
- Stripe Transfer Rate (STR):** 307200 KB/sec (DSW * DIOR)
- Seconds Between Tape Marks (SBTM):** 45
- Blocks Between Tape Marks (BBTM):** 27000 ((SBTM * DIOR * 1024) / MBS)
- Seconds Between Unbuffered Tape Marks (SBUTM):** 45
- Blocks Between Unbuffered Tape Marks (BBUTM):** 27000 ((SBUTM * DIOR * 1024) / MBS)

RAIT Options

- Parity Stripe Width:** 0 (0 = RAIT is disabled)
- Minimum Write Parity:** 0
- Number of RAIT Engines:** 0
- ☐ Read Verification
- ☐ Mount Minimum PVs

Buttons: Add, Start Over, Clone (partial), Clone (full), Dismiss

Setting field: Media Type

This window is used to manage tape storage class configurations.

Field descriptions.

Storage Class ID

The numeric identifier assigned to the storage class.

Storage Class Name

The descriptive name of the storage class.

Storage Class Type

The type of the storage class (always `Tape`).

Migration Policy

The migration policy associated with this storage class, or `None` if no migration is desired.

**Advice**

Do not configure a migration policy for a storage class at the lowest level in a hierarchy.

*If a migration policy is added to a storage class after files are created in the storage class, those files may never be migrated. Use the **mkmprec** utility to correct this problem. See the **mkmprec** man page for more information.*

Warning Threshold

A threshold for space used in this storage class expressed as a number of empty tape volumes. Alarms will be sent to SSM periodically when the number of empty tapes in the storage class falls below this value, and the **Space Thresholds** field on the *HPSS Health and Status* window will be changed to `Warning`. Note that this field will not have any effect if overridden by storage subsystem-specific thresholds.

Critical Threshold

Another threshold for space used in this storage class expressed as a number of empty tape volumes. Alarms will be sent to SSM periodically when the number of empty tapes in the storage class falls below this value, and the **Space Thresholds** field on the *HPSS Health and Status* window will be changed to `Critical`. Note that this field will not have any effect if overridden by storage subsystem-specific thresholds.

Optimum Access Size

The optimal transmission size to be used for a transfer request using this storage class. (Not currently used by HPSS. May be used by site-specific applications.)

Average Latency

The average time (in seconds) that elapses when a data transfer request is scheduled and the time the data transfer begins.

Maximum VVs to Write

The number of tape virtual volumes in the storage class that a Core Server will use for concurrent writes.

**Advice**

Small values in this field restrict files being written in the storage class to a small number of tapes, reducing the number of tape mounts. The number of tape drives used

to write files in the storage class will be limited to approximately the value of this field times the total stripe width of the mounted VVs. Read operations are not limited by this value.

Storage Segment Size

Media Type

The media type associated with the storage class.

Block Protection

When this flag is set, Logical Block Protect (LBP) will be enabled for tape volumes created in this storage class. When LBP is enabled, HPSS will generate and pass a CRC with each block written to the tape volume and validate the CRC accompanying each block read from the tape volume. This function is supported for the following media types.

- LTO (Generation 5 and greater)
- IBM 3592 (E07 and greater)
- Oracle T10000 (C and greater)

This option is not supported by virtual tape devices such as AWS Tape Gateway and should be disabled.

Media Block Size (MBS)

The Media Block Size is the size in bytes of a physical data block on the media. For tape, this can be almost any value within reasonable limits. If the tape hardware has a recommended physical block size, use that value.



Advice

*The **Media Block Size** should be set to a value appropriate for the volume type. See the Media block size selection section of the HPSS Installation Guide (in the HPSS planning chapter) for some recommendations.*

See also Device I/O Rate and Seconds Between Tape Marks in this section.

VV Block Size (VVBS)

The virtual volume block size is the number of bytes written to an element of a striped VV before the data stream moves to the next stripe element. It can be thought of as the stride length of striped data transfer operations. The length of the VV block has an effect on the striping efficiency. Short VV blocks cause more protocol overhead when writing striped devices. In non-striped applications, **VV Block Size** has little meaning so any convenient multiple of the **Media Block Size** will do.



Advice

When choosing a VV Block Size, the administrator should consider the characteristics of any data source or sink the storage class that will be copied to or from. Best performance of striped copies usually occurs when the VV Block Sizes of the source and sink are equal. This minimizes the data movement protocol overhead and helps to keep the data streams flowing smoothly.

VV Block Size must meet the following constraining requirements:

- *It must be an integer multiple of the Media Block Size.*
- *The PV Section Length (**Media Block Size (MBS)** multiplied by the **Blocks Between Tape Marks (BBTM)**) divided by the **VV Block Size (VVBS)** must be a whole number. For example, if the **Media Block Size (MBS)** is 64 KB, and the **Blocks Between Tape Marks (BBTM)** is 512, the physical volume section length is 32 MB. The **VV Block Size (VVBS)** could be 64 KB, 128 KB, 256 KB, or larger, but not 192 KB.*

See the Virtual volume block size selection (disk) and the Virtual volume block size selection (tape) sections of the HPSS Installation Guide (in the HPSS planning chapter) for information.

Data Stripe Width (DSW)

The number of physical volumes in a virtual volume in this storage class. In the case of RAIT VVs, this is the number of PVs that contain file data.

Data Stripe Length

The **Data Stripe Length** is the product of the **VV Block Size (VVBS)** and the **Data Stripe Width (DSW)**.

PV Size (PVSIZE)

The estimated size of the tape volume in bytes. The actual number of bytes that can be written to a tape will vary with individual media and data compressibility. Think of this field as the nominal length of a tape of this type.

VV SIZE (VVSIZE)

The virtual volume size. The product of the **PV Size (PVSIZE)** and the **Data Stripe Width (DSW)**.

Max Storage Segment Size

The upper bound for storage segment sizes created on volumes in this storage class. A value of "0" indicates that the maximum segment size is unlimited. If the value entered is nonzero, then it must be between 10% and 50% of the VV estimated size (**VV Size (VVSIZE)**).

Transfer Rate

Device I/O Rate (DIOR)

The approximate data transfer speed, in kilobytes per second (KB/sec), which can be achieved by devices corresponding to Media Type.



Advice

*This field is used in calculating a reasonable setting for **Seconds Between Tape Marks (SBTM)**.*

Stripe Transfer Rate (STR)

The estimated aggregate transfer rate for volumes of this type. This field is calculated as the **Data Stripe Width (DSW)** multiplied by the **Device I/O Rate (DIOR)**.

Seconds Between Tape Marks (SBTM)

The number of seconds between tape mark writes while writing continuously to tape.



Advice

*The **Device I/O Rate (DIOR)**, the **Media Block Size (MBS)** and the **Seconds Between Tape Marks (SBTM)** determine the length of a physical volume section. The administrator should set the **Device I/O Rate (DIOR)** to an appropriate value for the device, and then pick a **Media Block Size (MBS)**. Finally, adjust the **Seconds Between Tape Marks (SBTM)** to strike a balance between tape utilization and transfer efficiency.*

*Smaller values of **Seconds Between Tape Marks (SBTM)** cause more tape to be consumed in tape marks and break up data streaming throughput on streaming tape drives but cause controller buffers to be flushed more frequently and may improve access times to the middle of large files. Larger values improve transfer efficiency. However, smaller values may reduce the amount of data that may need to be rewritten on another tape when a tape reaches EOM. As tape drive streaming rates increase, we generally recommend larger values for this setting.*

Blocks Between Tape Marks (BBTM)

The maximum number of data blocks (of **Media Block Size (MBS)**) that will be written on a tape between consecutive tape marks.

The number of media data blocks between tape marks is calculated by SSM when the **Device I/O Rate (DIOR)** and **Seconds Between Tape Marks (SBTM)** are set. This value is displayed for information only.

The recording space between tape marks is called a "section". The PV Section Length is **Blocks Between Tape Marks (BBTM)** times the **Media Block Size (MBS)**. The VV Section Length is the PV Section Length multiplied by the **Data Stripe Width (DSW)**.

Seconds Between Unbuffered Tape Marks (SBUTM)

The number of seconds between synchronizing tape mark writes.



Advice

*When this value differs from **Seconds Between Tape Marks (SBTM)**, the Tape Mover is instructed to write buffered tape marks. A buffered tape mark will be written every **SBTM** seconds (or **BBTM** blocks), and an unbuffered/synchronizing tape mark will be written every **SBUTM** seconds (or **BBUTM** blocks).*

When buffered tape marks are enabled, write throughput is improved, but at the cost of potentially needing to rewrite more data in EOM cases. This is due to the behavior that when buffered tape marks are used, on an EOM or I/O error, the tape drive rolls the data written to tape back to the last tape mark that successfully made it to tape.

Blocks Between Unbuffered Tape Marks (BBUTM)

The maximum number of data blocks (of **Media Block Size (MBS)**) that will be written on a tape between synchronizing tape marks.

RAIT Options

Parity Stripe Width

When this value is greater than zero, the VV is a RAIT VV. This is the number of PVs in a VV that will carry parity information.

Minimum Write Parity

The minimum number of PVs that may carry parity. This value sets a limit on the number of PVs that may become unwritable in a volume before the VV becomes unwritable. For example, if a RAIT VV is configured with four data PVs and two parity PVs ($DSW = 4$, $PSW = 2$), then Minimum Write Parity may be set to "2" or "1". If the value is "2", then if one of the PVs becomes unwritable, the VV may no longer be written. If MWP is set to "1", then if one PV becomes unwritable, the VV remains writable, but if two become unwritable, the VV becomes unwritable.

Minimum Write Parity may also be set to zero. In that case any data written to the VV after all of the parity PVs have become unwritable is not protected with parity. Any additional losses of PVs may render the data written this way unrecoverable.

Number of RAIT Engines

The number of RAIT Engines that will be used to read and write VVs in this storage class. The system has the ability to use multiple RAIT Engines in a striped I/O fashion to increase throughput.



Advice

*Unless it can be shown that the RAIT Engine is a bottleneck in a particular configuration, set **Number of RAIT Engines** to one.*

Read Verification

When selected, data written on VVs in the storage class will be verified using the parity information stored with the data. When not selected, the RAIT Engine reads the data from the VV but does not check it using the parity information.

Mount Minimum PVs

VVs created with this option will be read using the smallest number of PVs necessary to recover the data. If **Read Verification** is off, the number of PVs mounted will be equal to the value of DSW. If **Read Verification** is on, the number of PVs mounted will be equal to $(DSW + 1)$.

6.1.4. Storage class subsystem thresholds

These windows allow the user to manage storage class space thresholds which are specific to a storage subsystem. These windows are accessible only from the *Configured Storage Classes* window. The user must first select a storage class from the **Storage Classes List**. The selected storage class will be referred to as the "parent" storage class to the subsystem threshold window which is opened when the storage class is selected.

Each storage class configuration window has fields for **Warning Threshold** and **Critical Threshold**. By default, these thresholds apply to all storage subsystems using the storage class. However, the

storage class thresholds can be overridden by configuring a subsystem-specific threshold allowing a storage subsystem to have its own customized set of **Warning Threshold** and **Critical Threshold** values.

After creating or modifying any thresholds, the changes will take effect after clicking the **Apply Config** button on the *HPSS Health and Status* screen.

Related information.

Section 4.2.3.8, “Storage class threshold overrides”

6.1.4.1. Disk Storage Subsystem-Specific Thresholds

Disk Storage Subsystem-Specific Thresholds

File Edit Help

Storage Class ID: 1

Storage Class Type: Disk

Storage Class Name: Storage Class 1

Default Warning Threshold: 80 percent

Default Critical Threshold: 90 percent

Storage Subsystem Count: 3

Subsys ID	Subsys Name	Warning	Critical
1	Subsystem #1	default	default
2	Subsystem #2	default	default
3	Subsystem #3	default	default

3/3

Subsys ID: 1

Subsys Name: Subsystem #1

Warning: default percent

Critical: default percent

Update Set To Defaults

Dismiss

This window is used to define warning and critical thresholds unique to a particular storage subsystem, overriding the values defined in the disk storage class. The user may modify either the **Warning** percent, **Critical** percent or both for one or more of the listed subsystems. Select a subsystem from the list and then modify the values on the lower portion of the window. When the new values have been entered, select **Update** to commit the changes. To remove the overridden threshold values, select the **Set To Defaults** button. The changes will be committed and displayed in the **Subsystem List** table.

After any modifications to the subsystem-specific thresholds, the changes will take effect after clicking the **Apply Config** button on the *HPSS Health and Status* screen.

Field descriptions.

Storage Class ID

The ID of the storage class to which these subsystem-specific thresholds apply.

Storage Class Type

The type of storage class (disk or tape).

Storage Class Name

The descriptive name of the storage class.

Default Warning Threshold

The default warning threshold for this storage class expressed as a percentage of used space.

Default Critical Threshold

The default critical threshold for this storage class expressed as a percentage of used space.

Storage Subsystem Count

The number of storage subsystems in the HPSS system.

Subsystem List

A list of configured storage subsystems. Columns in the list include:

Subsys ID

The ID of the selected storage subsystem.

Subsys Name

The name of the selected storage subsystem.

Warning

The current warning threshold value for the selected subsystem. If the storage class defaults are to be used, the text `default` will be displayed.

Critical

The current critical threshold value for the selected subsystem. If the storage class defaults are to be used, the text `default` will be displayed.

Buttons.

Set To Defaults

A button to remove the threshold override values from the selected storage subsystem. Select a storage subsystem row from the **Subsystem List**. The selected storage subsystem information will be displayed in the lower panel. Clicking the **Set To Defaults** button will commit the changes to the selected subsystem and `default` will be displayed for both the warning and critical thresholds.

Related Information.

Section 4.2.3.8, “Storage class threshold overrides”

6.1.4.2. Tape Storage Subsystem-Specific Thresholds

Tape Storage Subsystem-Specific Thresholds

File Edit Help

Storage Class ID: 32

Storage Class Type: Tape

Storage Class Name: Storage Class 11

Default Warning Threshold: 10 volumes

Default Critical Threshold: 5 volumes

Storage Subsystem Count: 3

Subsys ID	Subsys Name	Warning	Critical
1	Subsystem #1	default	default
2	Subsystem #2	default	default
3	Subsystem #3	default	default

3/3

Subsys ID: 1

Subsys Name: Subsystem #1

Warning: default volumes

Critical: default volumes

Update Set To Defaults

Dismiss

This window is used to define warning and critical thresholds unique to a particular storage subsystem, overriding the values defined in the tape storage class. The user may modify either the number of **Warning** volumes, **Critical** volumes or both for one or more of the listed subsystems. Select a subsystem from the list and then modify the values on the lower portion of the window. When the new values have been entered, select **Update** to commit the changes. To remove the customized threshold values, select the desired subsystem and click the **Set To Defaults** button. The changes will be committed and displayed in the **Subsystem List** table.

After any modifications to the thresholds, the changes will take effect after clicking the **Apply Config** button on the *HPSS Health and Status* screen.

Field descriptions.

Storage Class ID

The ID of the storage class to which these subsystem-specific thresholds apply.

Storage Class Type

The type of storage class (disk or tape).

Storage Class Name

The descriptive name of the storage class.

Default Warning Threshold

The default warning threshold for this storage class expressed as a number of free volumes.

Default Critical Threshold

The default critical threshold for this storage class expressed as a number of free volumes.

Storage Subsystem Count

The number of storage subsystems in the HPSS system.

Subsystem List

A list of configured storage subsystems. Columns in the list include:

Subsys ID

The ID of the selected storage subsystem.

Subsys Name

The name of the selected storage subsystem.

Warning

The current warning threshold value for the selected subsystem. If the storage class defaults are to be used, the text `default` will be displayed.

Critical

The current critical threshold value for the selected subsystem. If the storage class defaults are to be used, the text `default` will be displayed.

Buttons.**Set To Defaults**

A button to remove the threshold override values from the selected storage subsystem. Select a storage subsystem row from the **Subsystem List**. The selected storage subsystem information will be displayed in the lower panel. Clicking the **Set To Defaults** button will commit the changes to the selected subsystem and `default` will be displayed for both the warning and critical thresholds.

Related information.

Section 4.2.3.8, “Storage class threshold overrides”

6.1.5. Changing a storage class definition



When changing a storage class’s configuration, the Core and Migration/Purge Servers must be reinitialized or restarted in order for the changes to take effect. It is recommended to use the **Apply Config** button on the *HPSS Health and Status* window to reinitialize the servers.

If you’re changing a storage class’s definition as part of migrating to a new generation of storage media technology, refer to *Section 8.7, “Storage technology replacement”* for details and a high level procedure.

A storage class definition can be changed by bringing up the *Storage Class Configuration* window and making the desired modifications. Most fields in an existing storage class configuration are directly editable. The few that aren't include **Data Stripe Length**, **VV Size**, **Stripe Transfer Rate**, and **Blocks Between Tape Marks**. When planning to alter various values in a storage class's configuration, take the following aspects into account:

Migration Policy

If files have been stored in a storage class without a migration policy, and a migration policy is subsequently configured for it, the files created before the addition of the policy will not be migrated. Use the **mkmprec** utility to create migration records for these files so that they will migrate properly. See the **mkmprec** man page for more information.

Migration or Purge Policies

When a migration or purge policy is modified, or when a storage class configuration is modified to use a different migration policy, the affected Migration/Purge Servers and Core Servers must be reinitialized or restarted in order for the new policy to take effect. It is recommended to use the **Apply Config** button on the *HPSS Health and Status* window to reinitialize the servers.

Optimum Access Size, Device I/O Rate, Average Latency

These values are used only in an advisory capacity by HPSS. A customer may write a special interface using the Client API that takes advantage of the COS Hints capability and, in so doing, uses these values. Generally, they can be changed at will, but should be set to accurately reflect their intended usage and the backend storage devices serving the storage class.

Maximum VVs To Write

Change this as desired.


Max Multiplier, Min Multiplier, Average Number of Storage Segments

These values can be changed as desired for *disk* storage classes. They do not apply to *tape* storage classes.

Media Type

Care must be taken when changing a storage class's **Media Type**.

Changing this value does not apply to existent HPSS virtual volumes (VVs); a VV's storage class configuration is set upon resource creation and remains immutable for the lifetime of the VV. A new **Media Type** will only apply to new VVs that are created after modifying the storage class.

After selecting a new **Media Type**, an icon () will appear next to various values in the *Storage Segment Size* and *Transfer Rate* sections that automatically update to match the newly selected **Media Type**. It's important to be attentive to the values that change. Closely vet them against the values that were in the storage class configuration prior to changing the **Media Type**. The values can change in unexpected and undesired ways. For example, an LTO-4 storage class's 32 MB **VV Block Size** may revert to 1 MB upon changing the storage class's **Media Type** to LTO-6. The new 1 MB **VV Block Size** is not likely to be the desired final value. See the *Storage characteristics considerations* section of the *HPSS Installation Guide* (in the *HPSS planning* chapter) for guidance on configuring these storage class values.

Data Stripe Width and Parity Stripe Width

Care must be taken when changing a storage class's **Data Stripe Width** and **Parity Stripe Width**. File Hashing is not supported for striped storage classes, including RAIT storage classes.

Therefore, if the storage class is modified to be striped (by changing **Data Stripe Width** to be greater than 1), then the storage class must not be used in any Class of Service which supports File Hashing. This means that the **Hash Algorithm** in any COS configuration which references a hierarchy which includes this storage class must be set to "None." Logical Block Protection (LBP) is not supported for RAIT. If the storage class is modified to support RAIT (by changing **Parity Stripe Width** to a non-zero value), then **Block Protection** must be turned off and the related fields in the **RAIT Options** section must be configured.

Changing this value does not apply to existing HPSS virtual volumes (VVs); a VV's storage class configuration is set upon resource creation and remains immutable for the lifetime of the VV. A new **Data Stripe Width** or **Parity Stripe Width** will apply only to new VVs that are created after modifying the storage class.

6.1.6. Deleting a storage class definition

Before deleting a storage class definition, be sure that all of the subsystem-specific warning and critical thresholds are set to `default`. If this is not done, one or more threshold records will remain in metadata and will become orphaned when the storage class definition is deleted.

To delete a storage class definition, ensure that no files exist in this storage class and it is no longer referenced in any hierarchy. All of the virtual volumes in this storage class should also be deleted, using the Delete Resources or Reclaim operation. Once the storage class is empty and it is not referenced from any hierarchy definition, the storage class definition can be removed.

See also the **dump_acct_sum** utility, which can quickly verify the number of files in each COS.

It is recommended you call HPSS support for assistance with this operation.

6.2. Storage hierarchies

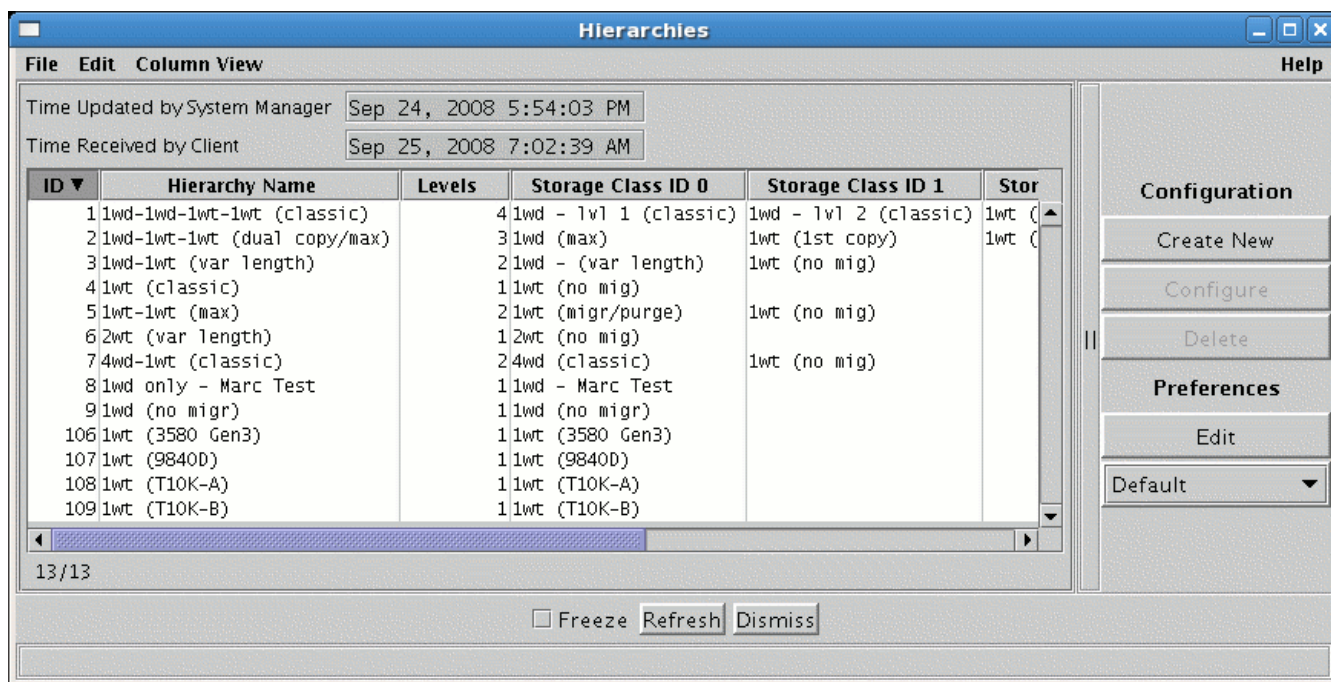
This section describes how to create, modify, and delete storage hierarchies.

6.2.1. Hierarchies window

An HPSS storage hierarchy can be created and managed using the *Hierarchies* window. This window is accessed from the *HPSS Health and Status* window's **Configure** menu, submenu **Storage Space**, item **Hierarchies**. Refer to *Section 3.9.3, "HPSS Health and Status"*.

The following rules for creating storage hierarchies are enforced by the *Hierarchies* window:

- A storage class may be used only once per hierarchy.
- Disk migration may migrate to one or more target levels in the hierarchy. To create multiple copies of data, select multiple migration targets.
- Tape migration may only migrate to a single target level in the hierarchy.
- A given storage class may be a migration target for only one storage class above it in the hierarchy (two different storage classes in the hierarchy may not both migrate to the same storage class).



This window allows you to manage storage hierarchies. It lists all currently defined hierarchies and allows you to change existing hierarchies. New hierarchies can also be created from this window.



When a new hierarchy is created or updated, the Migration/Purge Servers and the Core Servers must be reinitialized or restarted in order for the addition to take effect. It is recommended to use the **Apply Config** button on the *HPSS Health and Status* window to reinitialize the servers.

Field descriptions.

Hierarchy list

ID

The unique ID assigned to the storage hierarchy when it was configured.

Hierarchy Name

The descriptive name of the hierarchy.

Levels

The number of storage classes that have been configured into the hierarchy.

Storage Class ID 0 , Storage Class ID 1 , Storage Class ID 2 , Storage Class ID 3 , Storage Class ID 4

The name of the storage class at the given hierarchy level. HPSS arranges the storage classes into levels internally.

Configuration buttons.

Create New

Open a *Storage Hierarchy* window with default values.

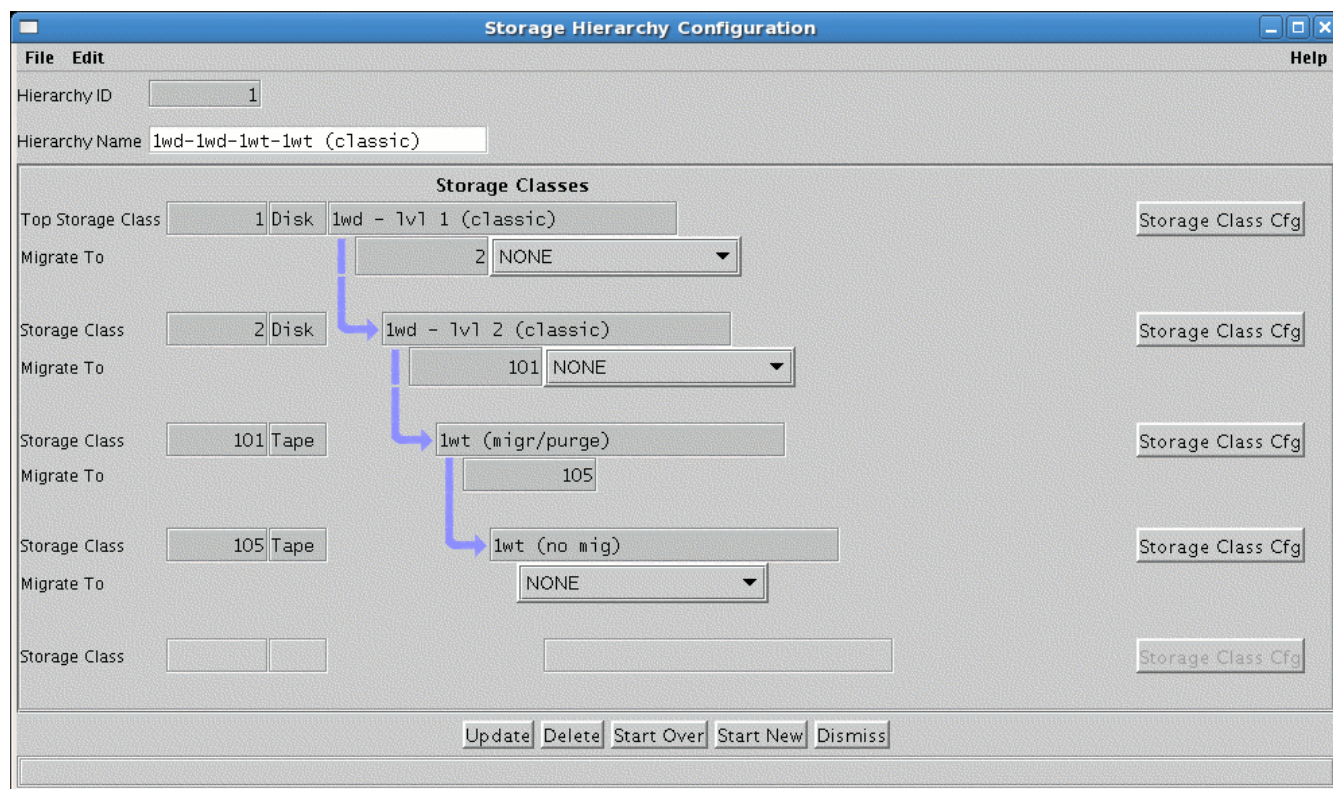
Configure

Open the selected storage hierarchy configuration for editing. One hierarchy from the list must be selected before this button is active.

Delete

Delete the one or more selected storage hierarchies.

6.2.2. Storage Hierarchy Configuration window



This window allows an administrator to manage a storage hierarchy.

A maximum of five storage classes can be configured into a hierarchy. A single storage class may appear in multiple hierarchies. *No hierarchy can reference the same storage class more than once.*

Note that in addition to defining the hierarchy, the appropriate migration policies must be configured for each storage class in the hierarchy. Without the proper migration policies, neither migration nor duplicate tape copies will be maintained. See *Section 6.4, “Migration policies”*.

Once a storage hierarchy has been created, it should not be deleted as long as there is a Class of Service entry associated with it. A Class of Service entry should never be deleted unless all files in that Class of Service have been deleted and no client is still attempting to use that Class of Service.

When a hierarchy is created or updated, the Core and Migration/Purge Servers must be reinitialized or restarted in order for the changes to take effect. It is recommended to use the **Apply Config** button on the *HPSS Health and Status* window to reinitialize the servers.

Field descriptions.

Hierarchy ID

The ID associated with this hierarchy. Any unique, positive 32-bit integer value. The default value is the last configured ID plus 1.

Hierarchy Name

The descriptive name associated with this hierarchy. The default value is "Hierarchy <ID>".

Top Storage Class

The storage class at the highest level of the hierarchy.

After a storage class is selected, a new storage class drop-down list will appear in the **Migrate To** field. Selecting a storage class from the drop-down list will define the storage class to which the **Top Storage Class** will migrate, and an arrow will show the migration path from the **Top Storage Class** to the second storage class.

This process may be repeated to create more migration paths.

Depending on whether the **Top Storage Class** is disk or tape, either one or two new storage class drop-down lists will appear.

Migrate To

The one or more storage class IDs to which the storage class will migrate. Initially, when creating a new hierarchy, these fields will not be visible. As the user configures the hierarchy, the **MigrateTo** values will be displayed. After a storage class has been selected from a drop-down list, the drop-down list will be replaced with a non-editable storage class ID corresponding to the user's selection. One or more new storage class drop-down lists will become visible.

Storage Class

The ID number of the storage class at the second through the fifth level of the hierarchy. See description above for **Top Storage Class**.

Storage Class Cfg button

Clicking this button will open the *Storage Class Configuration* window corresponding to the storage class name to the left of the button.

6.2.3. Changing a storage hierarchy definition

A storage hierarchy definition can be changed by bringing up the *Storage Hierarchy Configuration* window and making the desired modifications. Note that the only modification that is possible to an existing storage hierarchy is the inclusion of additional storage levels at the bottom of the hierarchy. All other modifications require that the hierarchy be deleted and recreated.

When adding a migration target to the bottom of a storage hierarchy, be sure that the migrating storage class is configured with appropriate migration and purge policies. Without migration and purge policies, no migration and purge will take place even if target levels exist in the hierarchy. Furthermore, any pre-existing data in the storage class (data which was written into the storage class before the storage class was configured to use migration and purge policies) will not be migrated and purged even after the migration and purge policies are added. The migration and purge policies are honored only for data written after the policies are put into effect. To migrate and purge data written before these policies became effective, use the **mkmprec** utility.

6.2.4. Deleting a storage hierarchy definition

The rules that apply for deleting a Class of Service also apply for deleting a storage hierarchy because in most systems a one-to-one relationship exists between hierarchies and Classes of Service. Before

deleting a hierarchy, all the files in all of the Classes of Service that use the hierarchy must be deleted. Failure to do this will render the files unreadable.

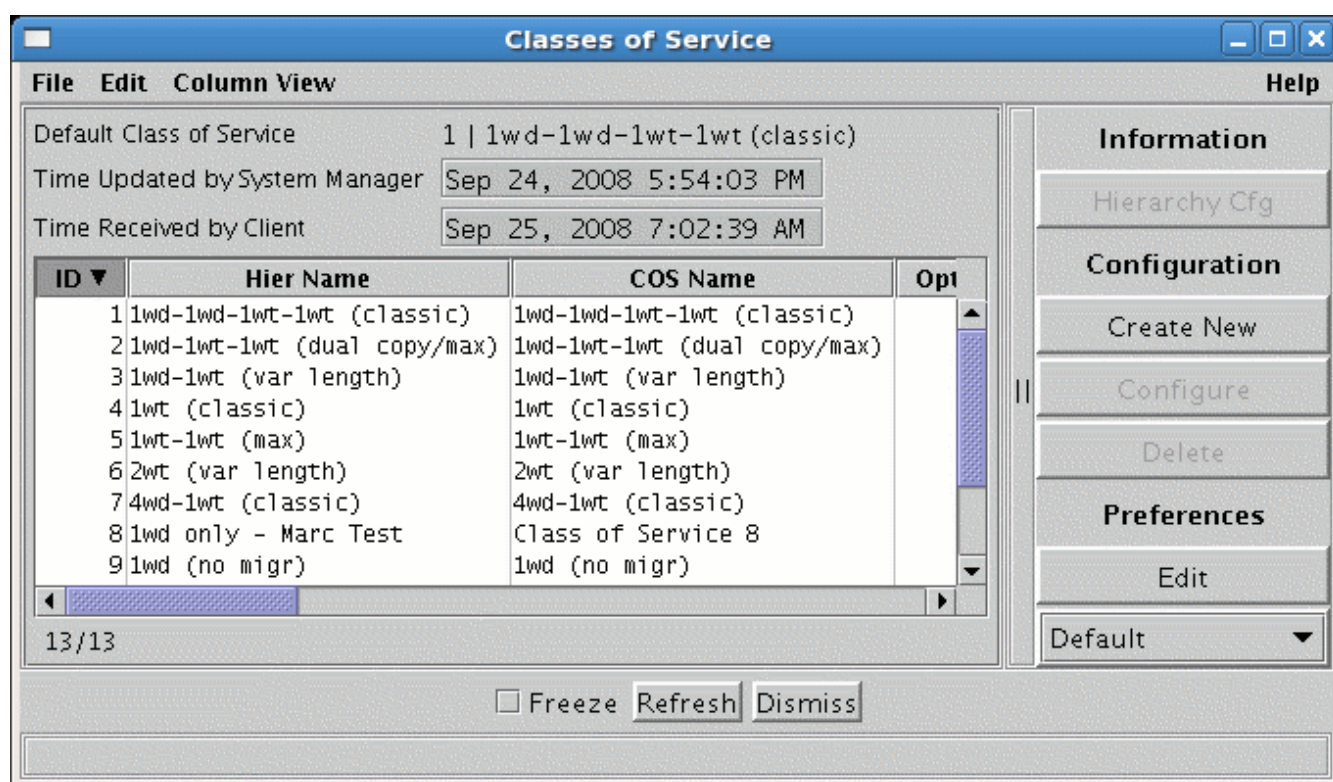
It is recommended you call HPSS support to assist with this operation.

6.3. Classes of Service

This section describes how to configure Classes of Service.

6.3.1. Classes of Service window

A COS can be created and managed using the *Classes of Service* window.



This window lists the Classes of Service that are currently configured. It also allows an administrator to update and delete existing Classes of Service and to add a new Class of Service. This window is displayed by selecting the **Configure** menu from the *HPSS Health and Status* window, then selecting the **Storage Space** submenu, and then selecting the **Classes of Service** submenu item.



When a new Class of Service is created or updated, the Migration/Purge Servers and the Core Servers must be reinitialized or restarted in order for the addition to take effect. It is recommended to use the **Apply Config** button on the *HPSS Health and Status* window to reinitialize the servers.

Field description.

Default Class of Service

The default Class of Service. This is displayed for information only and may not be changed from this window. It may be changed from the *Global Configuration* window.

Class of Service List

The central panel in this window shows the list of configured Classes of Service. The columns display the fields from the *Class of Service Configuration* window discussed in the following section.

Information buttons.

Hierarchy Cfg

Open a *Storage Hierarchy Configuration* window for the hierarchy associated with the selected Class of Service.

Configuration buttons.

Create New

Open a *Class of Service* window containing default values for a new Class of Service.

Configure

Open the selected Classes of Service configurations for editing.

Delete

Delete the selected Classes of Service.

6.3.2. Class of Service Configuration window

The screenshot shows a window titled "Class of Service Configuration (on hogwarts6)". The window has a menu bar with "File", "Edit", and "Help". The main area contains several configuration fields and sections:

- Class ID:** 2001
- Class Name:** Small Disk 1w -> T10KD 1w
- Storage Hierarchy:** 2001 | Small Disk 1w -> T10KD 1w
- Stage Code:** On Open
- Minimum File Size:** 0 bytes
- Maximum File Size:** 65535 bytes
- Allocation Method:** Fixed Length, Classic Style
- Hash Algorithm:** None
- Migration Order:** Directory
- Class Flags:**
 - ☒ Enforce Maximum File Size
 - ☒ Force Selection
 - ☒ Auto Stage Retry
 - ☒ Auto Read Retry
 - ☐ Auto Change COS
 - ☒ Truncate Final Segment
 - ☒ Full Aggregate Recall
- Class Characteristics:**
 - Access Frequency:** Daily
 - Optimum Access Size:** 4194304 bytes
 - Average Latency:** 0 seconds
 - Transfer Rate:** 4096 KB/sec
- R/W Operations:**
 - ☒ Read
 - ☒ Write
 - ☒ Append

At the bottom of the window, there are buttons for "Update", "Delete", "Start Over", "Clone (partial)", "Clone (full)", and "Dismiss".

This window allows an administrator to manage a Class of Service.

Field descriptions.**Class ID**

The unique integer ID for the COS. Any positive 32-bit integer value.

Class Name

The descriptive name of the COS. A character string up to 31 bytes in length. The default value is "Class of Service <ID>".

Storage Hierarchy

The name of the storage hierarchy associated with this COS.

Stage Code

A code that indicates the file staging options. Valid values are `On Open`, `On Open Async`, `No Stage`, and `On Open Background`. For the first COS created in a new HPSS system, the default value is `On Open`. For all subsequently created COSs, the default value is the same as the most recent COS configured.

**Advice**

Changing the Stage Code should be done with care. See Section 6.3.3, “Changing a Class of Service definition” for detailed information on each of the choices for Stage Code.

Minimum File Size

The size, in bytes, of the smallest bitfile supported by this COS. Valid values are any positive 64-bit integer value.

Maximum File Size

The size, in bytes, of the largest bitfile supported by this COS. Valid values are any positive 64-bit integer value.

Allocation Method

How disk storage segments will be allocated in the storage class at the top level of the hierarchy used by this COS. Ignored if the top level storage class is tape. Three allocation methods are available:

Fixed Length, Classic Style

All storage segments allocated for a single file, except possibly the last segment of the file, are the same size. This size is computed at file creation time to be the size which wastes the least disk space, based on the file size, the minimum and maximum segment sizes configured for the storage class, and the average number of segments per file configured for the storage class. In order to take maximum advantage of this allocation method, the application creating the file must specify the file size in its file creation hints.

The application can force this fixed segment size to be the maximum segment size configured for the storage class by specifying `HINTS_FORCE_MAX_SSEG` in the file creation hints.

If the application does not force the segment size to be the max size and it does not provide hints telling the Core Server the file size, then the Core Server will use the minimum segment size configured for the storage class.

If the file is purged completely from the top level and later restaged to the top level, a new segment size will be computed at stage time based on the current file size.

Fixed Length, Max Style

The size of each storage segment allocated for a file, except possibly the last segment, is the maximum segment size configured for the storage class.

Variable Length

The size of the first storage segment allocated for a file is the minimum segment size configured for the storage class. The size of each successive storage segment is twice that of the preceding segment, up to the maximum segment size configured for the storage class. Once the max segment size is reached, each remaining segment except possibly the last is also the max segment size.

Under all three allocation methods, the last segment of the file may be truncated to a size to best fit the actual data contained in the segment. See the **Truncate Final Segment** option in the **Class Flags** section below.

Hash Algorithm

This is the default hash algorithm to be used when generating digests for files in this COS as they are migrated or copied to tape, or written directly to tape. The following hash algorithms are supported.

- SHA1
- SHA224
- SHA256
- SHA384
- SHA512
- MD5
- CRC32
- ADLER32

This functionality is supported on storage hierarchies that contain a disk storage class at the initial level, with that storage class migrating to one or more tape storage classes. Because of the sequential nature of the supported hash algorithms, digest generation is currently limited to data being migrated or copied to non-striped tape volumes.

Another limitation involves the ability of HPSS to have the most current data spread over multiple levels in the storage hierarchy. This most often occurs when a file is appended after a duration long enough for the previous data to be migrated and purged from the initial disk storage level. If this occurs the new data will eventually be migrated, but the original data would already be valid at a lower level and will not be migrated. In this case, the file digest will not be generated or verified and a warning message will be issued.

This functionality is also supported when writing directly to tape. The class of service must contain a single stripe tape as the top tier storage class, and the file must be written from

beginning to end without any backseeks or parallel write operations. Concurrent access to the file will invalidate the hashing process. Subsequent writes to the file after the initial checksum has been generated will invalidate the file hash. The file hash will be finalized when the file is closed and compared against any user hash or stored as a generated checksum.

A user-provided digest can be associated with a file and a system generated digest can be obtained using programming interfaces.

Migration Order

The migration ordering criteria for this COS. The following ordering is supported:

Create Time (Default)

Migration will be ordered by file create time. This is the default.

Directory

Migration will be ordered by directory. Each directory's files eligible for migration will be grouped together. The depth of the directory grouping is one deep. Thus, files in a sub-directory will be grouped with other files in the sub-directory and not with files in the parent directory.

The objective is to co-locate data on tape, with either the directory of the files or the files' create times being the co-location criteria. A site may want to combine this setting with *COS Class Flag Full Aggregate Recall* to potentially minimize the number of tape operations required to bring a directory's files back from tape. Of course, the migration policy will continue to honor the migration policy settings for when a file is eligible for migration and how often migration will run. Refer to the *Section 6.4.2.1, "Disk Migration Policy configuration"* section for further disk migration options.

When modifying the **Migration Order**, the System Manager will notify the Migration/Purge Server about the change. The change won't affect in-flight migrations, but subsequent migrations will use the new **Migration Order** setting.

Class Flags

Enforce Maximum File Size

When set, a bitfile larger than the **Maximum File Size** cannot be created in this COS.

Force Selection

A flag to determine how a COS will be selected. If ON, a client must explicitly select this COS in order to have a file assigned to it; if the client merely supplies general COS hints for a file, this COS will not be selected.

Auto Stage Retry

When this flag is turned on, and a valid secondary copy of the data exists, and a stage from the primary copy fails, HPSS will automatically retry the stage using the secondary copy.

Auto Read Retry

When this flag is turned on, and a valid secondary copy of the data exists, and an attempt to read the first copy fails, HPSS will automatically retry the read using the secondary copy. This retry will not be attempted if any data was transferred between the client and HPSS during the failure.



*For **Auto Stage Retry** and **Auto Read Retry** to work properly, the COS must contain at least two copies of the files and at least one valid second copy must have been created during HPSS migration processing. If any tape storage class in the COS has a migration policy, that policy must be configured for **Migrate Files** or **Migrate Files and Purge**.*

Auto Change COS

When this flag is turned on, upon disk migration, files within this Class of Service (COS) that fall outside the **Minimum File Size** and **Maximum File Size** will have their COS changed to a target COS that has the appropriate file size range. The Migration/Purge Server initiates the change in Class of Service and the Core Server performs it as a background task. COSes with the **Force Selection** flag on will not be considered for the target COS. If multiple COSes meet the criteria, then the Migration/Purge Server will select the first COS that meets the criteria; if no COSes meet the criteria, then the file will remain in the current COS and be migrated according to the current COS's hierarchy. If the Core Server's **COS Copy to disk** flag, described in *Section 5.2.2, "Core Server-specific configuration"*, is not on and there is a tape level in the target COS's hierarchy, then the Core Server will not automatically copy the file to the disk level and will instead copy to the tape level. To take advantage of migration policies for the new COS's hierarchy (like tape aggregation), it is recommended that you turn **COS Copy to disk** on.

Truncate Final Segment

A flag to influence whether the final segment of files will be truncated at file close time in order to save disk space.

When a file is closed, its final segment may be truncated to the smallest segment size which is valid for the storage class and will still hold the actual data written to the segment. This size will be a power of two multiple of the minimum segment size configured for the storage class. The segment will later be expanded to its original size if the file is reopened and new data is written to the segment or to a point in the file beyond the segment. Truncation of the final segment is performed only on the file data at the top level of the storage hierarchy and only if the storage class at the top level is disk.

Truncation of the final segment is controlled by the NOTRUNC_FINAL_SEG flag in the bitfile descriptor; if this flag is off, truncation is performed at file close, and if it is on, truncation is not performed.

The COS **Truncate Final Segment** flag influences whether the NOTRUNC_FINAL_SEG flag in the bitfile descriptor is set at file creation time. The user may specify a creation hint to turn truncation off even if the COS allows truncation, but truncation can not be turned on if the COS prohibits it. If the COS **Truncate Final Segment** flag is on, then at file creation time the NOTRUNC_FINAL_SEG flag will be set to off, unless the user specifies the NOTRUNC_FINAL_SEG hint. If the COS **Truncate Final Segment** flag is off, then the NOTRUNC_FINAL_SEG flag will be set to on at file creation time, and the user may not override this.

Users may also turn on the NOTRUNC_FINAL_SEG flag in their own files after file creation by use of the `hpss_SetFileNotrunc` Client API function, but they may not turn off the flag. Authorized callers may use the `hpss_SetFileNotrunc` function to turn the flag on or off for any file. Authorized callers are those such as `hpsssm` who have `CONTROL` permission on the Core Server's Client Interface ACL.



Advice

Truncation of the final segment normally saves disk space. However, if a file is frequently written, closed, reopened, and appended to, the repeated truncation and re-expansion of the final segment could result in fragmentation of the disk, as the space for the expansion is not guaranteed to be adjacent to the original segment. In addition, it is possible that the only extent available for the expansion is larger than needed, so some space could actually be wasted.

*If it is known that the users of a COS do frequent appends to their files, it is better to turn the COS **Truncate Final Segment** flag off to avoid these potential fragmentation and space waste issues. If no users do frequent appends, or if those who do can be relied upon to turn truncation off for their own files, or if the system administrator can easily identify files which are frequently appended and can turn off truncation on them individually, then the site might want to take advantage of the space savings for the remaining files and leave **Truncate Final Segment** on in the COS definition.*

One additional consideration is that truncating the final segment incurs a small performance penalty. In development testing, this penalty was measured at 14% for writing 100 files with truncation in comparison to writing them without truncation; this value may vary under different site configurations.

Full Aggregate Recall

When this flag is turned on, and a stage is requested for a file which resides within a tape aggregate, then all of the eligible files within the aggregate are staged. Ineligible files are those which have been deleted and those which already have data staged at a higher level.

Class characteristics

Access Frequency

The frequency, on average, for accessing files in this COS. Valid values are Hourly, Daily, Weekly, Monthly, Or Archive.

Optimum Access Size

The suggested number of bytes that should be written at one time for maximum efficiency. (Not currently used by HPSS; may be used by site-specific applications.)

Average Latency

The average time, in seconds, that elapses between the time a transfer request is accepted for processing and the time the data transfer begins.

Transfer Rate

The average throughput (in KB per second) that can be transferred using this COS.

R/W operations

Read

If ON, files in this COS are readable, subject to the permissions on each individual file.

Write

If ON, files in this COS are writable, subject to the permissions on each individual file.

Append

If ON, files in this COS can have new data appended to them, subject to the permissions on each individual file. This check box may only be ON if the **Write** check box is also ON. If the **Write** check box is OFF, the **Append** check box is disabled.

In all cases, whether an individual file can be read or written will be determined by the permissions on that particular file. It may be the case that a COS is RW, but certain files within that COS may have their permissions set such that they can be neither read nor written.

6.3.3. Changing a Class of Service definition

A COS definition can be changed by bringing up the *Class of Service Configuration* window and making the desired modifications. Fields on the *Class of Service Configuration* window that can be changed without major impact are the **Access Frequency**, **Optimum Access Size**, **Average Latency**, and **Transfer Rate**. These fields are advisory in nature and are currently used only by the Client API. Users of custom HPSS interfaces will have to assess the impact of any change.

The **Maximum File Size** can be changed, but care should be exercised when doing so. Increasing this value may result in storing files that are inappropriate for the hierarchy supporting this COS. For example, suppose a hierarchy is defined with disk at the top level with a storage segment size of 64 KB and the **Maximum File Size** in this COS is set to 1 MB. Changing the maximum file size to 2 GB in this class would be inappropriate because 32,768 storage segments (each 64 KB in size) would be needed to contain such a file.

Associated with the **Maximum File Size** is the **Enforce Maximum File Size** flag. Changing this can have a significant impact. Turning the flag on constrains files that are already larger than the **Maximum File Size** to their current size. Existing smaller files will be constrained to the **Maximum File Size**.

Changing **Minimum File Size** can have an impact on COS selection. Currently, the PFTP and FTP interfaces use the **Minimum File Size** to select an appropriate COS based on file size.

Changing the **Stage Code** should be done with care:

- The **On Open** option stages files to the top of the hierarchy when the file is opened, and is synchronous. The open operation returns when the stage operation completes.
- The **No Stage** option can have a significant impact on system performance. Repeated reads of files that have been migrated will generally be satisfied from lower levels in the hierarchy where files are likely to be on tape. If users are writing only parts of a file at a time with significant delays between the writes, migration may result in the file being stored in a more fragmented manner on tape.
- The **On Open Async** option causes stage operations to be started when files are opened, but open operations do not wait for stages to complete; they return to their callers immediately. If a caller then reads one of these files, the read operation will block until the stage is complete.
- The **On Open Background** option causes an open operation to complete immediately, returning a special error code that indicates that the stage is not complete. The caller can then poll HPSS for completion of the stage operation before reading the file. Normally this polling is done automatically by the Client API, but Client API polling can be turned off which will allow the client to do the polling.

Changing the **Auto Stage Retry** flag has no adverse side effects. It is recommended that this flag be enabled for all COSs using multiple copies.

The hierarchy associated with a COS cannot be changed without deleting all files in this COS or moving them to another COS. Failure to observe this condition will likely result in lost and corrupted files. Any changes to an existing COS other than changing the hierarchy ID can be put into effect by clicking the **Apply Config** button on the *HPSS Health and Status* window.

When modifying the **Migration Order**, the System Manager will notify the Migration/Purge Server about the change. The change won't affect in-flight migrations, but subsequent migrations will use the new **Migration Order** setting.

6.3.4. Deleting a Class of Service definition

To delete a COS, you must ensure that all files in the COS have been either deleted or moved to another COS and that all configuration references to the COS have been removed. These configurations include the Global Configuration's Default Class of Service, Storage Subsystem's Default COS Override, and every fileset's associated Class of Service.

Deleting a Class of Service is an unusual administrative action and HPSS does not provide utility programs to delete all the files in a given COS or move all files in a given COS to another COS. A certain amount of site-specific ad hoc planning and usage of available tools will be necessary to identify, remove, or move all files from a Class of Service. The **dump_acct_sum** utility program provides a count of the number of files in all Classes of Service based on the accounting metadata and was developed primarily so there would be a reasonable way to know if a Class of Service was empty.

To delete the COS definition, use the **Delete** button on the appropriate *Class of Service Configuration* window.

6.3.5. Changing a file's Class of Service

The Core Server provides a means to change the Class of Service of a file. The Core Server moves the body of the file as appropriate to media in the destination Class of Service, then allows the usual migration and purge algorithms for the new Class of Service to apply. The file body is removed from the media in the old Class of Service. If large numbers of files on tapes are changed, an opportunity to retire, repack, and reclaim tapes in the old Class of Service may arise.

The **scrub** utility can be used to initiate the COS change request of a file as follows:

```
% scrub
scrub> changecos <fullPathname> <newCOSid>
scrub> quit
```

COS changes are performed using standard HPSS internal copy mechanisms. These do not verify the file checksums during the copy. Sites should record the checksums of files to be COS-changed and verify them with **hpsssum** after the COS change is complete.

6.3.6. Canceling a Class of Service change request

A Class of Service change request can be canceled in order to deal with a persistent problem in changing the file.

The scrub utility can be used to cancel the COS change request as follows:

```
% scrub
scrub> changecos <fullPathname> 0
scrub> quit
```

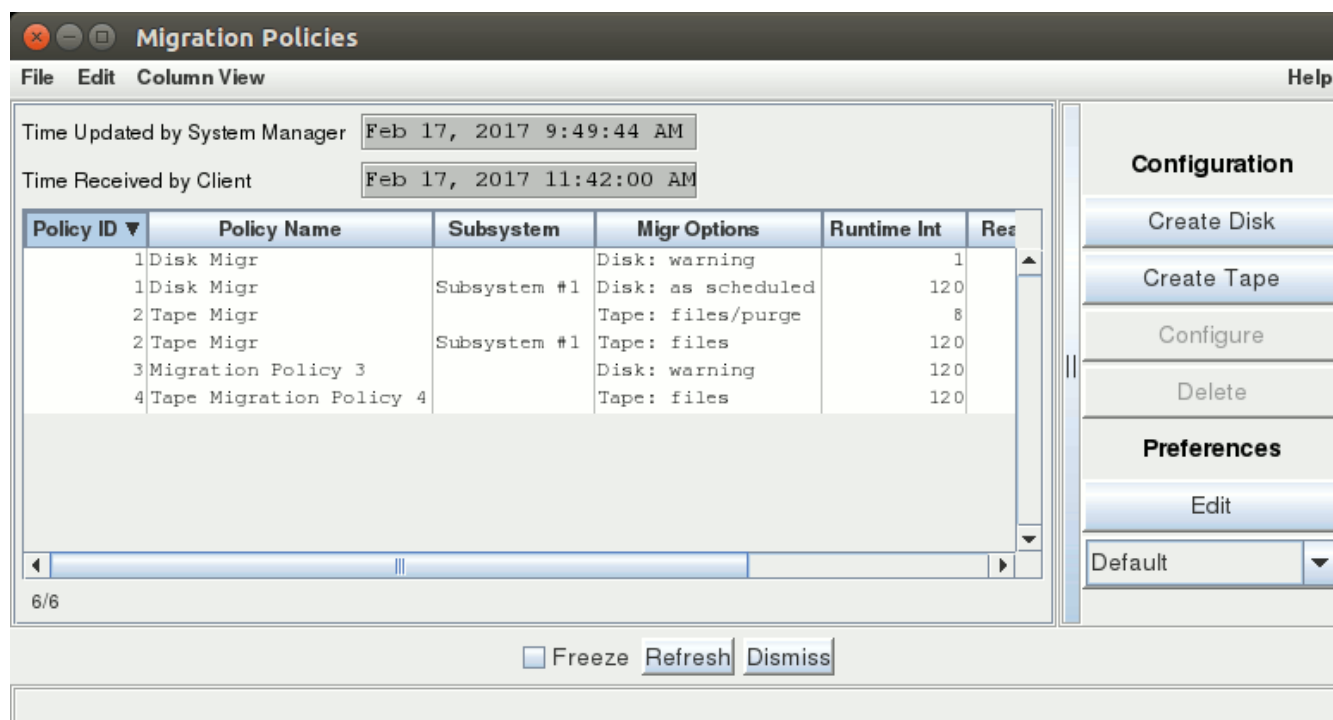
The associated Core Server COS change record will be removed immediately. However, this request may fail if the existing COS change request is currently being processed by one of the Core Server's COS change threads. In this case, the user should attempt to reissue the request later.

6.4. Migration policies

A migration policy is associated with a storage class and defines the criteria by which data is migrated from that storage class to storage classes at lower levels in the storage hierarchies. Note, however, that it is the storage hierarchy definitions, not the migration policy, which determine the number and location of the migration targets. Also, note that a storage class may be used in multiple hierarchies and may have different migration targets in each.

A basic migration policy must be assigned to any disk or tape storage class that requires migration services. This basic policy determines the migration parameters for the storage class across all storage subsystems. Storage subsystem-specific migration policies may be added which override the basic policy in the selected subsystems. The storage subsystem-specific migration policies share the migration Policy ID and Policy Name with the basic policy.

6.4.1. Migration Policies window



This window displays a list of all the configured migration policies.

Both basic and subsystem-specific migration policies are created and managed using the *Migration Policies* window. The basic policy must be created before creating any subsystem-specific policies.

The fields in the basic policy are displayed with default values. Change any fields to desired values as needed. Click on the **Add** button to write the new basic policy to the HPSS metadata.

To configure a subsystem-specific policy, select an existing basic policy and click on **Configure**. The basic policy will be displayed in a tabbed window with a tab for each storage subsystem for which subsystem-specific thresholds have been created. Click on the tab for the desired storage subsystem and make the necessary subsystem-specific changes, and then click on the **Update** button. To add a specific threshold configuration for a new subsystem, click the **New Subsystem Policy** button.



*When a migration policy is added to an existing storage class, the Migration/Purge Servers for that storage class and the Core Servers must be reinitialized or restarted in order for the policy to take effect. It is recommended to use the **Apply Config** button on the HPSS Health and Status window to reinitialize the servers.*

When HPSS needs to apply a migration policy to a subsystem, it first searches for a matching subsystem policy. If none is found, HPSS uses the basic policy. You can consider the basic policy to be the default for all subsystems which do not have policies of their own.

If a policy has one or more subsystem-specific overrides, the basic and subsystem-specific policies will be listed separately. Each entry will have the same ID, but the **Subsystem** field will be blank for the basic policy and contain the subsystem ID for each subsystem-specific override.

After changing any migration policy information, the changes will take effect after clicking the **Apply Config** button on the *HPSS Health and Status* screen.

Field descriptions.

Policy List

Policy ID

The unique ID number assigned to this migration policy.

Policy Name

The name assigned to this migration policy.

Subsystem

If the policy is a subsystem-specific override, this shows the subsystem to which it applies. This field will be blank for basic policies.

Other Migration Policy List columns

The remaining columns provide the same information that can be found in *Section 6.4.2.1, “Disk Migration Policy configuration”* and *Section 6.4.2.2, “Tape Migration Policy configuration”* windows.

Configuration buttons.

Create Disk

Opens a *Disk Migration Policy* window with default values.

Create Tape

Opens a *Tape Migration Policy* window with default values.

Configure

Opens the selected migration policy for editing.

Delete

Deletes the one or more selected migration policies.

Related information.

Section *Migration/Purge Server* of the *HPSS Installation Guide* (in the *HPSS planning* chapter)

6.4.2. Migration policy configuration

This section describes the configuration, update, and deletion of migration policies.

6.4.2.1. Disk Migration Policy configuration

The screenshot shows the 'Disk Migration Policy' configuration window. The 'Policy Name' is 'Migration Policy 1' and the 'Policy ID' is '1'. The 'Basic' tab is active. The configuration settings are as follows:

Setting	Value
Runtime Interval	120 minutes
Last Update Interval	60 minutes
Continue migration if greater than	0 minutes of I/O
Number of Migration Streams Per File Family	1
Total Migration Streams	1
Migrate when Warning Threshold Exceeded	<input checked="" type="radio"/>
Migrate when Critical Threshold Exceeded	<input type="radio"/>
Migrate only at Scheduled Intervals	<input type="radio"/>
Optimize for Even File Family Migration	<input type="checkbox"/>
Read-Optimized Aggregates	<input type="checkbox"/>
Max Aggregate Size	25TB

At the bottom of the window, there are buttons for 'Update', 'Delete', 'Start Over', 'New Subsystem Policy', 'Start New', and 'Dismiss'. A status bar at the bottom indicates 'Successfully saved the 'Disk Migration Policy''.

This window allows an administrator to manage disk migration policies and their subsystem-specific overrides.

Subsystem-specific policies define migration rules to be applied on a subsystem basis instead of using the default migration policy. When a migration run occurs and the storage class's migration policy does not have a subsystem-specific migration policy with the same subsystem ID, then the basic migration policy will be applied; otherwise, the matching subsystem-specific migration policy will be

used. Click the **Create Disk** button on the *Migration Policies* window to create a new policy. To bring up the configuration for an existing policy, select the policy from the list and click on the **Configure** button.

Field descriptions.

Policy Name

The descriptive name of the migration policy.

Policy ID

A unique ID associated with the migration policy.

Runtime Interval

The number of minutes to wait between the end of one migration run and the start of the next. This can be in the range [0 to 4,294,967,295].



Advice

*The **Runtime Interval** is the maximum number of minutes from the end of one migration run to the start of the next. Migration needs to run often enough so the storage class to which the policy is assigned does not fill up. Migrations can also be started manually from the Active HPSS Storage Classes window.*

Last Update Interval

The number of minutes that must pass since a file was last updated before it can become a candidate for migration.

Continue migration if greater than <#> minutes of I/O

This setting allows disk-to-tape migration to continue if it will result in I/O that runs for more than the specified amount of time; otherwise, if there's not enough data to meet the threshold, then the newly eligible files will be migrated on the next **Runtime Interval**, thus freeing up drives. For some customer workloads, this might be useful for reducing dismount/remount activity during migration. The overall idea is to avoid further tape mount activity unless enough data has accumulated to warrant the cost of mounting the tapes.

Consider, for example, an HPSS instance with multiple file families and a slow ingest rate for each family. The slow, consistent influx of files in each family causes migration to continue running, migrating just a few files at a time per family. This might be acceptable for customers that have drives dedicated to migration or are primarily concerned with getting files migrated as soon as possible. However, other customers might prefer for migration to start, migrate all the currently eligible files, and then only continue migrating data if some minimum amount of data has accumulated.

A value of "0" will be the default and will represent the default migration behavior: that disk-to-tape migration will continue as long as existent files are eligible for migration. A nonzero value will represent the minimum time for which migration I/O will need to run in order to continue migration.

This field will be ignored for disk-to-disk only hierarchies (that is, no tape).

To calculate the minutes of I/O, the MPS will look at the transfer rate of the storage classes in each level of the hierarchy involved in the migration; it will pick the slowest to use for the calculation.



Advice

Each file family migration will look for this field and determine whether or not it should continue. Thus if one family has enough data to continue, that family will continue to migrate; but if another family doesn't have enough, migration of that family's data will be delayed until the next runtime interval.

Verify that the Transfer Rate is correct in your storage classes.

Number of Migration Streams Per File Family

The number of migration streams to be allocated to each file family. This value effectively determines how many file families can be migrated simultaneously. For example: **Number of Migration Streams Per File Family** is set to "2" and **Total Migration Streams** is set to "10"; up to five families will be migrated simultaneously (five families with two streams working on each family yields a total of ten migration streams).

Total Migration Streams

The number of parallel migration threads which will be run during migration for each copy level in the hierarchy. This value must be a multiple of the *Number of Migration Streams Per File Family*.



Advice

*Setting the **Total Migration Streams** value too high may result in higher than desired tape drive utilization. Each hierarchy within a disk storage class is migrated in a separate but parallel fashion. As such, when there are files eligible for migration in each hierarchy within a disk storage class, tape mount requests will be generated for each hierarchy. The **Total Migration Streams** value may be used indirectly to control how many tape mounts are requested, and thus how many drives are used during migration. The maximum number of drives needed for a migration can be calculated as the total stripe width value (SW) multiplied by the **Total Migration Streams** value (TMS) for every hierarchy and disk-to-tape copy level within the storage class:*

$$(SW \times TMS)_{Hier1,LevelA} + (SW \times TMS)_{Hier1,LevelB} + \dots$$

$$(SW \times TMS)_{Hier2,LevelA} + (SW \times TMS)_{Hier2,LevelB} + \dots$$

+ ... +

$$(SW \times TMS)_{Hiern,Levelx} + (SW \times TMS)_{Hiern,Levely} + \dots$$

For example, an HPSS system that has a disk storage class which migrates to a tape storage class with total stripe width two. The migration policy's stream count is also two. In this situation, the maximum tape drive usage would be four: two sets of two-wide tape VVs.

A hierarchy that migrates data from disk to a first and second 4-wide tape copy with a stream count of one would use a maximum of eight tape drives: the first copy would go to one set of 4-wide tapes and the second copy would migrate to a second set of 4-wide tapes. A second hierarchy added to the mix with the same characteristics would double the number of tape drives needed.

Now consider an HPSS system that has disk storage class A at the top level of hierarchy 1 and hierarchy 2. Storage class A has a migration policy with a total stream count of 4. Hierarchy 1 migrates to a tape storage class with total stripe width of two. Hierarchy 2 migrates to a first copy tape storage class with total stripe width of two, and a second copy tape storage class with a stripe width of one. The total drive utilization for a maximal migration case (when there exist files eligible for migration in all hierarchies of the disk storage class) will be twenty:

$$(2_{SW} \times 4_{TMS})_{Hier1,Level1} + (2_{SW} \times 4_{TMS})_{Hier2,Level1} + (1_{SW} \times 4_{TMS})_{Hier2,Level2} = 20$$

*In addition to the **Total Migration Streams**, the **Maximum VVs to Write** parameter in the Tape Storage Class Configuration window may be used to further control how many tape drives/cartridges are used for migration: it will place a cap on the maximum number of VVs that can be written simultaneously in a storage class.*



See Migration Stream Distinction for additional information.

Triggered Migration Options

There are three choices for managing migration behavior when a storage class is running out of space and the next migration isn't yet scheduled to occur:

Migrate when Warning Threshold Exceeded

A migration run should be started immediately when the storage class warning threshold is exceeded.

Migrate when Critical Threshold Exceeded

A migration run should be started immediately when the storage class critical threshold is exceeded.

Migrate only at Scheduled Intervals

Do not trigger unscheduled migrations when storage class thresholds are exceeded.

Optimize for Even File Family Migration

Select this option in order to evenly migrate files from all file families with migratable files. Otherwise, the MPS will migrate all files within a file family before moving onward to the next file family.

Read-Optimized Aggregates

Select this option if you want aggregates optimized for read. **Read-Optimized Aggregates** will have only one ordering key value per aggregate (e.g. only files contained in Directory X). When **Read-Optimized Aggregated** is not selected, you will get Write-Optimized Aggregates which allows more than one ordering key value per aggregate (e.g. files contained in Directory X followed by files contained in Directory Y, etc). The default behavior is Write-Optimized Aggregates. This is another important aspect to consider when preparing to employ Full Aggregate Recall.

Max Aggregate Size

The maximum total size in bytes of a tape aggregate. Disk-to-tape migration bundles bitfiles into tape aggregates, and this value sets the maximum aggregate size. Tape aggregates are limited to

a maximum of the total estimated size of the destination tape volume. Tape aggregation can be disabled by setting the "Max Aggregate Size" to 0.



Advice

For details about migration of files to aggregates on tape, see Section 12.1, “Overview of tape aggregation”.

6.4.2.2. Tape Migration Policy configuration

This window allows an administrator to manage tape migration policies and their subsystem-specific overrides.

Subsystem-specific policies define migration rules to be applied on a subsystem basis instead of using the default (basic) migration policy. When a migration run occurs and the storage class's migration policy does not have a subsystem-specific migration policy with the same subsystem ID, then the basic migration policy will be applied; otherwise, the matching subsystem-specific migration policy will be used. Click the **Create Tape** button on the *Migration Policies* window to create a new policy. To bring up the configuration for an existing policy, select the policy from the list and click on the **Configure** button.

Field descriptions.

Policy Name

The descriptive name of the migration policy.

Policy ID

A unique ID associated with the migration policy.

Runtime Interval

The number of minutes to wait between the end of one migration run and the start of the next. This can be in the range [0 to 4,294,967,295].

Last Read Interval

The number of minutes that must pass since a file or tape volume was last read before it can become a candidate for migration.

Last Update Interval

The number of minutes that must pass since a file or tape volume was last written before it can become a candidate for migration.

Max Active File Delay

The maximum number of minutes beyond the larger of the Last Read and Last Update intervals to wait for the intervals to be satisfied before forcing the migration of files. This can be in the range [-1 to 2,147,483,647].

Example 6.1. Max Active File Delay values

A value of "-1" means wait to indefinitely for the intervals to be satisfied. A value of "0" means to wait no longer than the larger of the two intervals. A value of "2880" means to wait no more than two days beyond the larger of the two intervals. If a file or tape were to remain busy for two days beyond the initial Last Read and Last Update intervals, the MPS would force the migration of the delayed files even if the files or source tape volume were active.

Migration Target

The desired percentage of free tape space to have available when migration is completed. The migration will be terminated when this goal is reached or when no more files meet the migration criteria. This goal may not be attainable if the total size of all files not eligible for migration is large.

Total Migration Streams

This value determines the degree of parallelism in the file migration process. This applies to policies using the **Migrate Files** and **Migrate Files and Purge** options only (see **File Options** below).

File Options

There are two options available for determining how tape migration handles files. Only one of these options may be selected at a time.

Migrate Files

A duplicate copy of the tape files will be made at the next lower level in the hierarchy.

This enables a tape migration algorithm which is similar to disk migration. This algorithm is based on individual files rather than tape volumes, and is able to make second copies of files stored on tape. In this algorithm, individual files are selected for migration based on their last write time and the settings in the migration policy. The selected files are migrated downwards to the next level in the hierarchy. The order in which files are migrated is based approximately on the order in which they are written to tape. Prior to migration, however, the files are sorted by source tape volume, so their migration order may vary. In this algorithm, files at the source level are never purged. Note that the source tape storage class must be repacked and reclaimed manually.

Migrate Files and Purge

The tape files will be moved to the next lower level in the hierarchy. No duplicate copy is maintained.

This enables the same file-based tape migration algorithm as **Migrate Files**, but purges files from the source level when they are successfully migrated. Additionally, this algorithm considers both the read and write times for a candidate file against the settings in the migration policy. This algorithm is useful for moving inactive files to a lower hierarchy level. Note that the source tape storage class must be repacked and reclaimed manually. Also note that if migrated files are staged up in the hierarchy, HPSS does not run a purge policy to delete the cached copy at the higher level.

Related information.

Section 4.2.3.7, “Migration and purge policy overrides”.

6.4.2.3. Changing a migration policy

To update an existing basic migration policy, select the policy from the *Migration Policies* list window and click the **Configure** button. The disk or tape *Migration Policy* window will appear. Make the desired changes to the basic policy and subsystem policies by selecting the appropriate tabs and placing new values in fields. When all changes have been made, click the **Update** button to record the changes.

In order for changes made to a migration policy to take effect, the changes should be applied by clicking the **Apply Config** button on the *HPSS Health and Status* window. This will ensure the Migration/Purge servers and the Core servers are reinitialized and that the changes to the migration policy are applied to all storage classes that reference the policy. You can also reread the policy from the *MPS Storage Class Information* window or the **Migration Controls** selection list on the *Active Storage Classes* list screen. However, the changes are only applied to the storage class and storage subsystem for which the policy is reread.

6.4.2.4. Deleting a migration policy



*Before deleting a basic migration policy, make sure that it is not referenced in any storage class configuration. Users cannot delete a migration policy that is in use by a storage class. Attempts to delete a migration policy that is referenced will fail due to database referential integrity constraints. Be sure to check the **Migr Policy** column of the *Configured Storage Classes* list for any storage class that references the migration policy to be deleted.*

*Be sure to click the **Apply Config** button on the *HPSS Health and Status* window after deleting a migration policy.*

To delete a migration policy, select the policy from the *Migration Policies* list window and click the **Delete** button. If a basic policy is selected, and the policy has subsystem-specific policies associated with it, a prompt will appear asking if the basic policy and the related subsystem-specific policies should all be deleted since they must be deleted before the basic policy can be.

A migration policy may also be deleted by clicking the **Delete** button on the *Migration Policy* configuration window for the policy. Subsystem-specific policies can be individually deleted from the *Migration Policy* configuration window for the policy by selecting the specific subsystem's tab and clicking the **Delete Subsystem Policy** button.

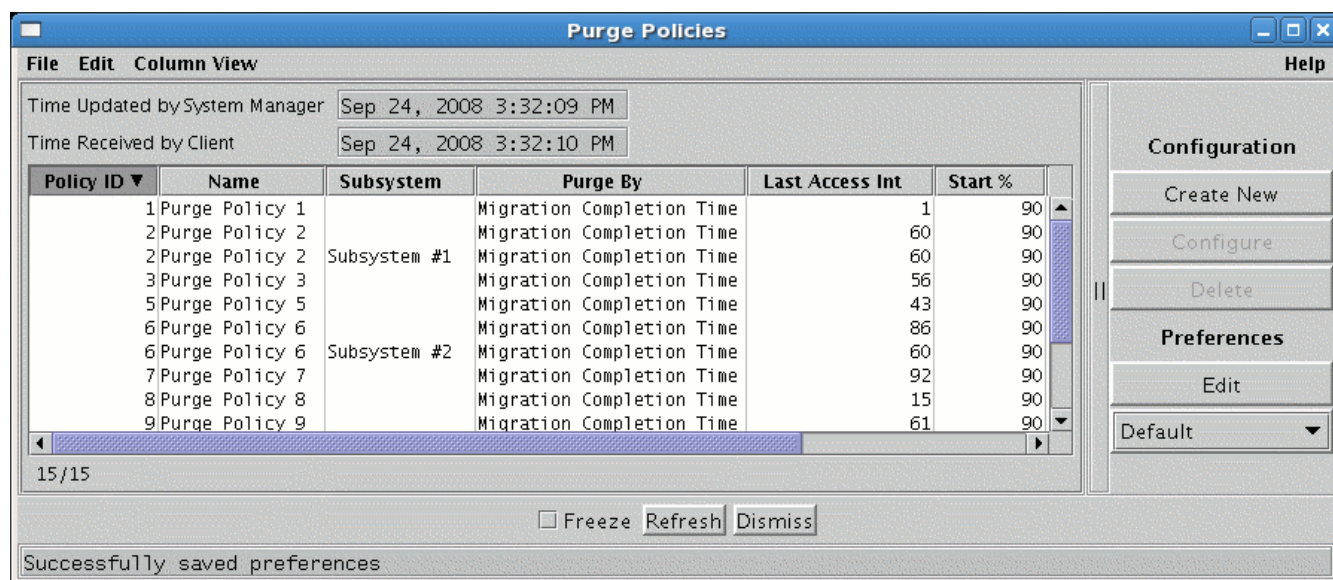
6.5. Purge policies

A purge policy is associated with a disk storage class and defines the criteria by which data is purged from that storage class once migration has copied that data to storage classes at lower levels in the storage hierarchies.

A basic purge policy must be assigned to all disk storage classes that require purge services. This basic policy determines the purge parameters for the storage class across all storage subsystems. If individual purge policies are desired for specific storage subsystems, storage subsystem-specific purge policies may be added which override the default values in the basic policy. The storage subsystem-specific purge policies share the same purge Policy ID and Policy Name with the basic policy.

Both basic and subsystem-specific purge policies are created and managed using the *Purge Policies* window. The basic policy must be created before creating the subsystem-specific policies.

6.5.1. Purge Policies window



This window displays a list of all the purge policies configured in the system.

To create a new purge policy, click the **Create New** button in the *Purge Policies* window. A new *Purge Policies* window appears. Enter the desired parameters for the policy, and then click the **Add** button. The policy will be added to the system. If subsystem-specific policies are desired, click the **New Subsystem Policy** button after the window refreshes, enter the specific purge policy parameters, and click the **Update** button. This process can be repeated for each subsystem.



*When a purge policy is added to an existing storage class, the Migration/Purge Servers for that storage class and the Core Servers must be reinitialized or restarted in order for the policy to take effect. It is recommended to use the **Apply Config** button on the HPSS Health and Status window to reinitialize the servers.*

Field descriptions.

Purge Policy List columns

The columns provide the same information that can be found on the *Purge Policies* window in the following section.

Configuration buttons.**Create New**

Open a *Purge Policies* window with default values for a new policy.

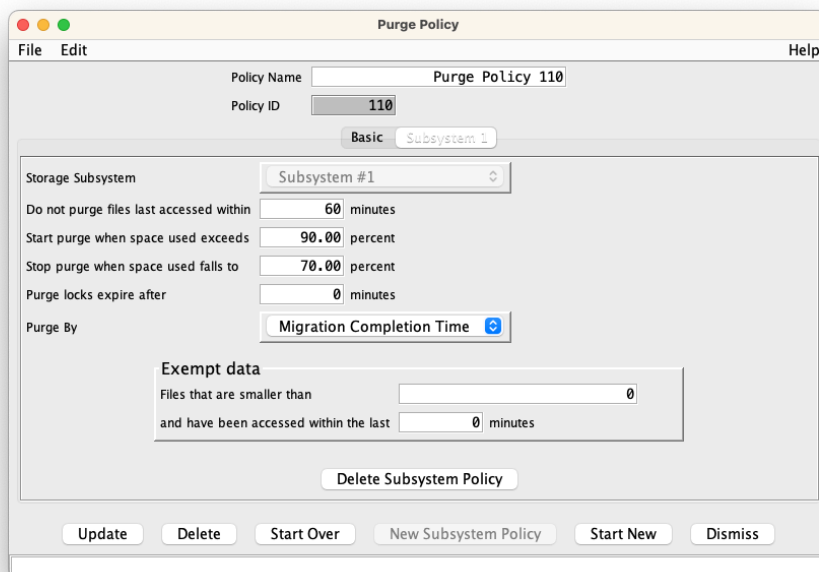
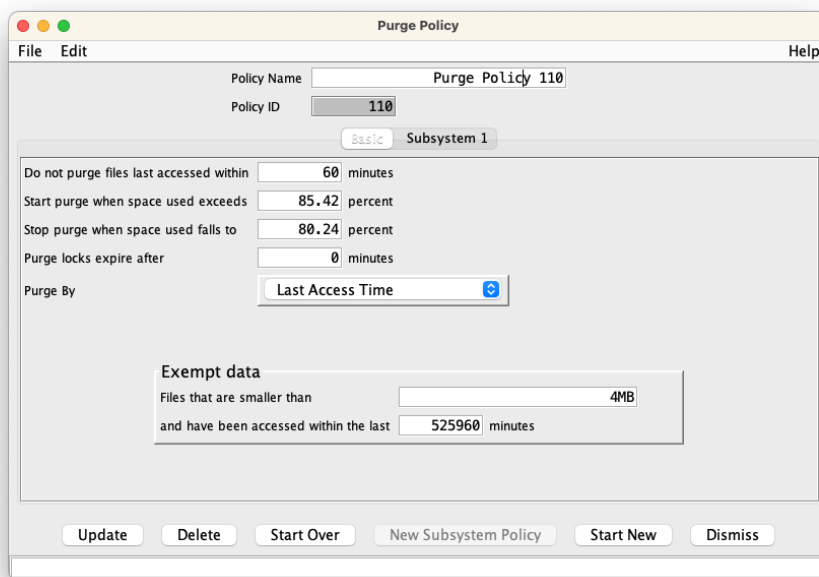
Configure

Open the one or more selected purge policies for editing.

Delete

Delete the one or more selected purge policies.

6.5.2. Purge Policy configuration



This window allows you to manage a purge policy. Purge policies are assigned to storage classes to tell the Migration/Purge Server and Core Server how to free disk space occupied by files which have been migrated. Purge policies apply to disk storage classes only.

The window always includes the **Basic** tab and may include one or more **Subsystem** tabs. These will be referred to as "basic" and "subsystem" below.

Each purge policy consists of a single basic policy and zero or more subsystem policies. Each subsystem policy has the same policy ID and name as its basic policy. When a new purge policy is needed, the basic policy must be created first; subsystem policies can then be created as needed.

When HPSS needs to apply a purge policy to a subsystem, it first searches for a matching subsystem policy. If none is found, HPSS uses the basic policy. You can consider the basic policy to be the default for all subsystems which do not have policies of their own.

After changing any purge policy information, the changes will take effect after clicking the **Apply Config** button on the *HPSS Health and Status* screen.

When creating a subsystem policy, fields which differ from the basic policy values can be thought of as exceptions to the basic policy values. However, this does not imply that unchanged defaults maintain any links with the basic policy. Changes in the basic policy have no effect on subsystem policies.

Field descriptions.

Policy Name

The descriptive name of the Purge Policy.

Policy ID

A unique ID associated with the Purge Policy.

Do not purge files last accessed within

A file will not be a candidate for purge until it has remained unaccessed (for read or write) for the length of time specified by this field.

Start purge when space used exceeds

Purge will begin for a storage class when the amount of its space used exceeds this threshold. Used space includes any file in the storage class, whether it has been migrated or not. This threshold may be expressed as a fractional percentage with a precision of up to two decimal places (for example, 85.42 percent).

Stop purge when space used falls to

Purging will stop for a storage class when the amount of its space used drops to this threshold. Note that the purge may stop before this point if it runs out of files which are available for purging. This threshold may be expressed as a fractional percentage with a precision of up to two decimal places (for example, 80.24 percent).

Purge Locks expire after

Maximum number of minutes that a file may hold a purge lock. Purge locked files are not eligible for purging. A value of "0" indicates that purge locks expire immediately.

Files may be purge locked to the highest level (level zero) of a hierarchy, provided that level zero is of type disk, and data exists at level zero. Once a file is purge locked, it is no longer a purge candidate. By entering a "lock expiration duration," you will be able to monitor the number of

purge locked files and the number of expired purge locked files using the Migration/Purge Server (MPS) log. After a purge run, the MPS will log the total number of purge locked files and the number of expired purge locks (those files that have been locked longer than this specified number of minutes). To view the names of the files with expired purge locks, use the Purge List Utility (**plu**) with the **-x** parameter on the given storage subsystem.

Super purge locks are not configurable by the "lock expiration duration." They do not expire and they are not yet reported by the MPS or the **plu** utility. Super purge locks should be rare. They are set temporarily by the **recover** utility and released when the **recover** utility determines that a safe copy of the file still exists. They are set internally by the Core Server whenever a manual force purge is performed on a file, which should not be done except with guidance from HPSS support. Super purge locks may be removed manually by the **scrub purgeunlock** command with the assistance of HPSS support.

Purge by

The MPS uses this time attribute of a file to determine which files are eligible to be purged.

By default, files are selected for purge based on the time the purge record was created, the Migration Completion Time. Alternately, the selection of files for purging may be based on the time the file was created, the File Creation Time, or the time the file was last accessed, the Last Access Time.



*If this field is changed, use the **Apply Config** button on the HPSS Health and Status screen to apply the change across the Core Servers and Migration/Purge Servers. Additionally, files will be purged in an unpredictable order until all purge records existing at the time of the change are cleared.*

Exempt data

The exempt data includes two conditions to prevent certain files from being identified as purge candidates. Both conditions must be true in order for a file to be exempt as a purge candidate. Otherwise, the file will be subject to the standard purge policy.

Files that are smaller than

A file that is the specified size or smaller in bytes may be ineligible as a purge candidate.

and have been accessed within the last

A file may not be a candidate for purge until the time elapsed for the most recent read or write time exceeds the value of this field.

Storage Subsystem (subsystem policy tab only)

The descriptive name of the storage subsystem to which a subsystem-specific policy applies. This field is filled in with the selected storage subsystem name at the time a subsystem-specific policy is created and may not be changed afterwards.

Related information.

Section 4.2.3.7, "Migration and purge policy overrides".

6.5.3. Changing a purge policy

To update an existing basic purge policy, select the policy from the *Purge Policies* window and click the **Configure** button. After modifying the basic policy, click on the **Update** button to save the

changes. To update an existing subsystem-specific purge policy, first select and display the existing basic policy and then select the tab for the desired subsystem policy. After modifying the subsystem-specific policy, click on the **Update** button to save the changes.

In order for changes made to a purge policy to take effect, the changes should be applied by clicking the **Apply Config** button on the *HPSS Health and Status* window. This will ensure the Migration/Purge servers and the Core servers are reinitialized and that the changes to the purge policy are applied to all storage classes that reference the policy. You can also reread the policy from the *MPS Storage Class Information* window or the **Purge Controls** selection list on the *Active Storage Classes* list screen. However, the changes are only applied to the storage class and storage subsystem for which the policy is reread.

6.5.4. Deleting a purge policy



Before deleting a basic purge policy, make sure that it is not referenced in any storage class configuration. Users cannot delete a purge policy that is in use by another storage class. Attempts to delete a purge policy that is referenced will fail due to database referential integrity constraints. Be sure to check the Purge Policy column of the Configured Storage Classes list for any storage class that references the purge policy to be deleted.

*Be sure to click the **Apply Config** button on the HPSS Health and Status window after deleting a purge policy.*

To delete a purge policy, select the policy from the *Purge Policies* window and click the **Delete** button. If a basic policy is selected, and the policy has subsystem-specific policies associated with it, a prompt will appear asking if the basic policy and the related subsystem-specific policies should all be deleted since the basic policy cannot be deleted until all of the subsystem-specific policies are deleted.

A purge policy may also be deleted by clicking the **Delete** button on the *Purge Policy* configuration window for the policy. Subsystem-specific policies can be individually deleted from the *Purge Policy* configuration window for the policy by selecting the specific subsystem's tab and clicking the **Delete Subsystem Policy** button.

6.5.5. Purge parallelism

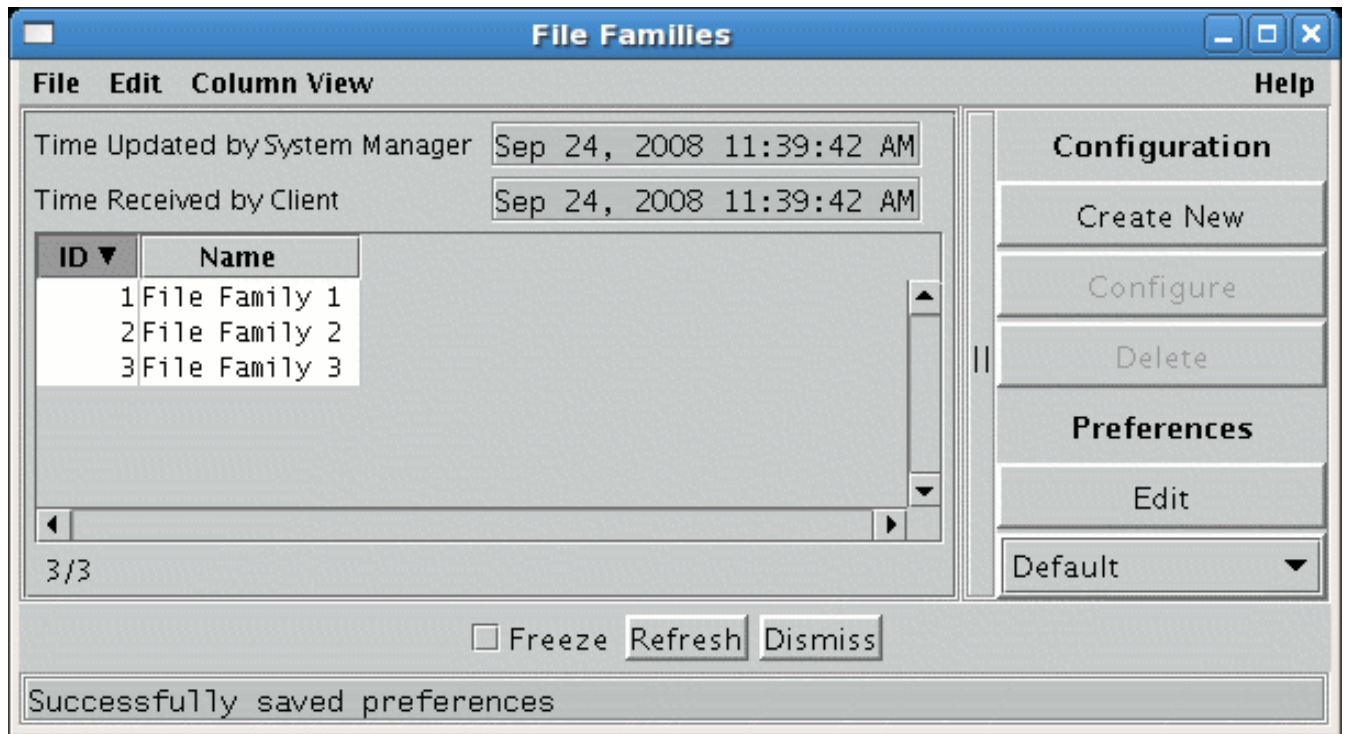
The degree of parallelism used for purging files from disk storage classes can be controlled via the HPSS environment. By default, the MPS degree of parallelism is set to two threads per storage class. You *may* find that increasing the number of purge threads increases the speed at which your HPSS system can purge files from disk. An increase in the number of purge threads will cause a slight increase in the amount of memory used by the MPS.

To run with a custom number of purge threads, use the HPSS environment variable named `HPSS_MPS_PURGE_PARALLELISM`. The valid range of values for this variable is [1,10]. As an example, for the MPS to run with four threads of purge per storage class, the following needs to be specified in `<HPSS_PATH_ETC>/env.conf`:

```
HPSS_MPS_PURGE_PARALLELISM=4
```

After changing the degree of purge parallelism, restart the MPS for it to begin using the new configuration.

6.6. File families



This window lists the configured file families.

Field descriptions.

The fields of the columns of this window are those of the file family configuration described in *Section 6.6.1, "File Family Configuration"*.

Configuration buttons.

Create New

Open a *File Family Configuration* window with default values.

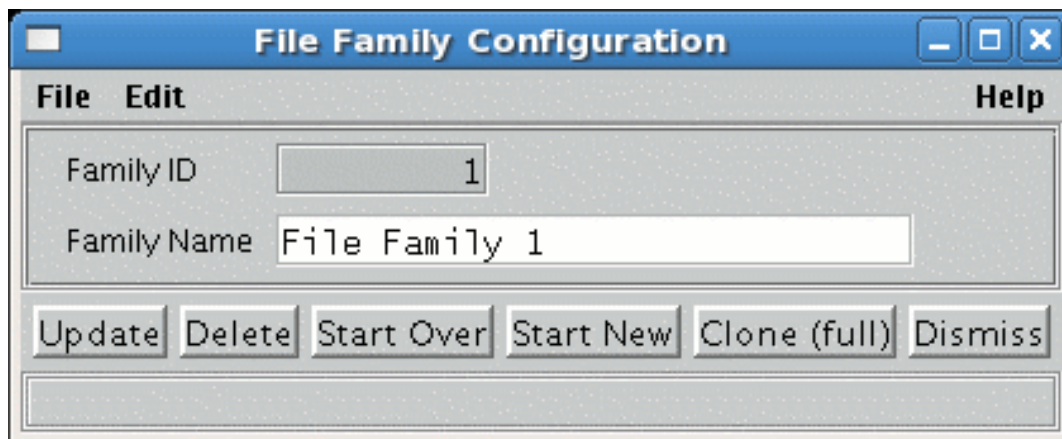
Configure

Open the selected file family configurations for editing.

Delete

Delete the selected file family configurations.

6.6.1. File Family Configuration



This window allows you to manage a file family configuration.

Field descriptions.

Family ID

A positive integer which serves as a unique identifier for this file family.

Family Name

The descriptive name for this file family. The name should be unique among all file families in this HPSS installation. The default value is "File Family <ID>".

6.6.2. Changing a file family

The only change that can be made to a file family is to change the **Family Name**. This action has no side effects and takes effect when the **Update** button is clicked.

6.6.3. Deleting a file family

SSM provides a means to delete a file family definition, but this action has no effect on files already recorded in the family. The file family attribute is permanent and cannot be removed or changed in existing files. Deleting a file family just means that files can no longer be added to the deleted file family.

Chapter 7. Device and drive management

Every disk and tape drive that is used by HPSS is controlled by two servers. The PVL performs mounts and dismounts (for disk devices these are logical operations only), and the Movers perform I/O. In support of these two views, the terms "PVL drive" and "Mover device" are used to refer to the configuration and information maintained about the drive by the PVL and Mover, respectively.

All configuration information is managed through a single SSM window. However, there are separate PVL drive and Mover device information windows, each with their own administrative state and other settings. When administrative changes are made, it is important to know whether they are to be performed against the PVL drive, the Mover device, or both.

Devices and drives are added and deleted dynamically. This means that they are added and deleted while the system is up. The PVL must be running and it is expected that the associated Mover and PVR (for tape) are also running. A new device and drive cannot be used until the PVL has successfully informed the associated Mover and PVR about the new device and drive. Likewise, a previously deleted device and drive cannot be added back until the PVL has successfully informed the associated Mover and PVR about the delete, and a previously added device and drive cannot be deleted until the PVL has successfully informed the associated Mover and PVR about the add. Additional details are described later in this chapter.

Devices and drives can be updated dynamically (that is, while the PVL, associated Mover, and associated PVR are running) via configuration update. The drive (tape) must be locked and idle before an update can occur. Additionally for disk drives, the volume must be unallocated before an update can occur. However this restriction can be omitted for some drive attributes for which the configuration metadata can be updated dynamically via the *PVL Drive Information* window. Similarly there are some device attributes for which the server "in memory" values can be changed dynamically (but not preserved when the Mover is restarted) via the *Mover Device Information* window. Additional details are described later in this chapter.

Note that all newly-created PVL drives are locked (that is, the *Devices and Drives* window **Drive Admin State** reports "Locked"), which prevents mounts from occurring on the drive. Therefore, these drives will need to be unlocked before they can be used. This can be done by highlighting the device/drive in the *Devices and Drives* window and clicking the "Unlock" button under the **Drive Administration** group heading.

Review the appropriate subsections of *Section 5.2.7, "Physical Volume Repository (PVR)-specific configuration"* before configuring a device/drive or modifying its state.

7.1. Configure a new device and drive

Be certain all tape drives are configured to use variable block sizes. See *Section 7.1.2, "Enable variable block sizes for tape devices"*. These steps are also covered in detail in the *Tape devices* subsection of the *HPSS Installation Guide* (within the *Mover* section of the *HPSS planning* chapter).

The Device and Drive configuration entry can be created and managed using the *Devices and Drives* window (see *Section 7.1.1, "Devices and Drives window"*).

After a drive has been configured, its state can be changed (unlocked versus locked) using the state change instructions described in *Section 7.1.1, “Devices and Drives window”*.

For disks using a file system interface, sparse files need to be created. The UNIX **truncate** utility is a good way to do this. The file name of the file will go in the **Device Name** field when configuring the device. Some file systems have features (such as software RAID or erasure coding) that can be used to enhance the integrity of the data stored on them, or they have features (such as compression) that increase the capacity of the storage media. These features can be used by the file system, but do not affect HPSS processing.

Before adding a tape drive to HPSS, the administrator must first fully configure the tape hardware and driver software into the host platform and operating system. For disk drives, the raw disk must first be defined in the operating system. Note that using the block special file name in the device configuration may result in a large performance degradation. For SAN disk devices, the **hpss_san3p_part** utility with the **-i** option and device name must be run to assign a UUID to the disk device. This UUID should be used in the **Device Name** field when configuring the disk device.

Currently the number of drives which may be configured per PVR is limited to 256. The current maximum number of PVRs is 64. The maximum number of devices per Mover is also 64. Given these limits, the total number of tape drives allowed in the system is 16,384.

Devices and drives are added dynamically. This means that they must be added while the system is running. The PVL must be running and it is expected that the associated Mover and associated PVR (for tape) are also running. A new device/drive cannot be used until the PVL has successfully notified the associated Mover and PVR (for tape) about the new device/drive. If the PVL is unable to successfully notify the Mover or PVR about the new device/drive, it will retry several times; each failure will be logged as an ALARM. After the last failed attempt to notify the Mover or PVR, the PVL will abandon retrying. The administrator is advised to stop and start the Mover (if it was up during the failed retries) and to reinitialize the associated PVR (if it was up during the failed retries) before the device/drive can be used.

Likewise, a previously deleted device/drive cannot be added back until the PVL has successfully notified the associated Mover and PVR (for tape) about the previous delete. If the administrator attempts to add a device/drive with a previously deleted Drive ID that is still pending delete notification, the administrator will get a BUSY failure and the pending notification will be aborted. An ALARM will be logged informing the administrator that the create failed since the Mover/PVR notification pends, and another ALARM will be logged informing the administrator that the previous drive delete notification will be aborted and to try later. Wait a minute or so before retrying in order to give the PVL time to abort the notification. The **Drive Flag** field in the *Devices and Drives* window and *PVL Drive Information* window will report whether or not a particular drive is pending Mover/PVR notification. Of course, this is only viewable while the drive exists.

At this point, the Administrative State of the drive must be unlocked before it is available. Refer to *Section 7.3.1, “Unlocking a drive”*.

Disk Device Configuration

Disk Device Configuration

File

Edit

Help

Device ID

5

Device Type

Generic - Default Disk

Mover

Mover (camille)

Basic Information

Device Name

/opt/disk/disk1

Media Block Size

4KB

Capacity

1TB

Starting Offset

0

SAN ID

0

Device Flags

☒ Read Enabled

☒ Write Enabled

☐ SAN3P-transfer Enabled

☒ Multiple Mover Tasks

☐ Disable Direct I/O

Drive Address

Controller ID

5

Maintenance Date

(ex. Dec 31, 2022 10:59:30.715 AM)

Mounted Volume

CA000100

Comment

Update

Delete

Start Over

Clone (partial)

Clone (full)

Dismiss

Tape Device Configuration

Tape Device Configuration

File Edit Help

Device ID: 24

Device Type: IBM - 3592 E07/EH7 (TS1140) Tape

Mover: Mover (camille)

Basic Information

Device Name: /dev/tape/by-id/scsi-3500507630f31680a

Device Flags

☒ Read Enabled
 ☒ NO-DELAY Support
 ☒ Removable Media Support
☒ Write Enabled
 ☐ TAOS Support
 ☒ Locate Support
☒ Fast Positioning
☐ Disable RAO

PVR: SCSI PVR E07

Drive Address: T10: IBM|03592E07 0000013D0372

Controller ID: 24

Polling Interval: -1 seconds (-1 to disable)

Maintenance Date: (ex. Dec 31, 2021 10:59:30.672 AM)

Comment:

Library Unit ID: -1

Update Delete Start Over Clone (partial) Clone (full) Dismiss

Update the selected device/drive: Drive must be locked.

These windows allow you to manage a tape or disk device/drive configuration.

Modifying/Updating the device/drive configuration via the *Tape Device Configuration* or *Disk Device Configuration* windows is managed dynamically by the PVL. The associated Movers, or the associated PVRs (for tape) are expected to be running. If you are changing the Mover and/or PVR, then both the current Mover/PVR and the new Mover/PVR will need to be running. An option exists which allows the PVL to perform the update without associated servers running. This option should be used with caution as it mainly addresses the abnormal condition where a Mover platform has failed and drives are being moved to a new platform/Mover and should be used with caution.

Additionally PVL drive configuration metadata updates can occur dynamically via the *PVL Drive Information* window. Likewise, some Mover device configuration "in memory" updates (that is, metadata not changed) can occur dynamically via the *Mover Device Information* window. Refer to Section 7.1.3, "Changing a drive's configuration".

Field descriptions.**Device ID**

The ID associated with this device/drive. Any positive 32-bit integer value.

Device Type

The type of device over which data will move.

Mover

The name of the Mover that controls the device.

**Advice**

A single Mover can manage a maximum of 64 devices. When the Mover is started and more than 64 devices are found, it will log an error and exit.

Device Name

The name by which the Mover can access the device. This name is usually the path name of a device special file such as `/dev/rmt0`. For SAN disks, this name is instead a UUID assigned to the disk device by running the **hpss_san3p_part** utility with the **-i** option and device name. The UUID for a SAN disk device can be determined by running the **hpss_san3p_part** utility with the **-s** option.

**Advice**

For locally attached disk devices, the pathname should refer to the raw/character special file (for example, `/dev/rhpss_disk1`).

For Linux systems, raw disk devices are deprecated so HPSS will support block devices and block device partitions. The SCSI disk block devices are referred to by pathnames of the form `/dev/sdXY`, where X begins at "a" and is incremented for each LUN detected and Y begins at "0" and is incremented for each partition of the block device. See `sfdisk(8)` for details about partitioning Linux block devices. When a block device is specified, the Mover will use `O_DIRECT` when opening the device and thus bypass the normal buffering supported by Linux block devices. SCSI tape devices are referred to by pathnames of the form `/dev/stX`, where X begins at zero and is incremented for each LUN detected.

*For SAN3P devices on Linux, Linux DM Multipath can be used instead of the default SCSI block devices. This is done by setting the `HPSS_SAN3P_LINUX_DMMPATH` environment to "on". The environment setting can be added with the **hpss_env_set** command, or by simply using your favorite editor to change the `/var/hpss/etc/env.conf` file. The setting should be added for each Mover and client that needs to access a SAN3P device. If you have previously been using a SAN3P configuration, deletion of the `/var/tmp/hpss_san3p.id.*` files on each client and Mover may be required. These files will be regenerated for the next SAN3P access.*

*For disk storage technologies that are using a file system interface, use the file name of the sparse file. In addition, when filling in the **Capacity** field, use the size of the sparse file.*

Media Block Size (disk only)

The block size for the device. This value should be a power of two multiple and a multiple of the underlying disk block size; otherwise, an error will occur at first I/O.

Capacity (disk only)

The size of the device (or sparse file) in bytes.

**Advice**

The storage class to which this drive will be assigned must have a PV Size less than or equal to this value.

As of version 9.2, the HPSS disk size limit has increased from 256 TiB to 1 EiB. If you plan to use a disk larger than 256 TiB, contact HPSS support for advice.

If the Starting Offset is nonzero (SAN3P disk only), then the Capacity value cannot be greater than the actual size of the underlying device less the Starting Offset value.

If this value is modified after the disk has been imported into HPSS, it must be emptied, exported, and re-imported.

Starting Offset (SAN3P disk only)

For default disks this field is always zero. For SAN3P disks this is the offset in bytes from the beginning of the disk logical volume at which the Mover will begin using the volume. The space preceding the offset will not be used by HPSS. This value should be a multiple of the **Media Block Size**.

**Advice**

This field is used only for SAN3p disk devices. The provided default value will specify the offset of the next available byte based on previous devices allocated using the same SAN disk UUID. If changing default value take great care that the new starting offset will not result in the device overlaying any other device using the same SAN disk.

SAN ID (disk only)

The ID for the SAN group. A group is a set of devices that are logically related, with the primary characteristic being that all devices in group are accessible by all clients. Each group is identified with this globally unique group ID. It is assigned to all devices within a SAN group.

Device Flags

The following fields are the device flags used by the Mover.

Read Enabled

An indication of whether the device is available for reading.

Write Enabled

An indication of whether the device is available for writing.

Locate Support (tape only)

An indication of whether the device supports a high speed (absolute) positioning operation.



Advice

This option is supported for 3590, 3590E, 3580, 3592, 9840, 9940, T10000 and GY-8240 devices.

NO-DELAY Support (tape only)

An indication of whether the device supports opening the device with no delay flag set, while allowing tape I/O operation after the open.



Advice

On some tape devices, this will allow for a quicker polling operation when no tape is presently loaded in the device. This field is meaningful for tape devices only.

Removable Media Support (tape only)

An indication of whether the device supports removable media.

TAOS Support (tape only)

An indication of whether the device is set up to use the Spectra Logic TAOS feature. TAOS enables an RAO-like interface for devices which do not support the RAO command set.



Advice

What is TAOS? [https://spectralogic.com/wp-content/uploads/white_paper_increasing_file_read_throughput_on_tape_with_taos.pdf]

TAOS Support should only be enabled for devices which have already been set up for TAOS with a Spectra Logic library. Discuss your data use cases with Spectra Logic and HPSS before enabling, as TAOS can result in a lower performance in certain scenarios. The TAOS device setup can be confirmed by using `device_scan` or `lsscsi` on the Mover device. Verify that Spectra Logic supports TAOS on the device type prior to enabling this flag.

Fast Positioning (tape only)

An indication of whether the device should skip section header verification when positioning via LBA.

Disable RAO (tape only)

An indication of whether this device should have RAO (Recommended Access Order) capability disabled. This field has no effect for devices which do not support RAO.

SAN3P-transfer Enabled (disk only)

If ON, SAN3P data transfers to this device will be supported.



Warning: There is a security vulnerability associated with the use of SAN3P. If a user is root on a machine which has access to the SAN (for example, a client machine) then that user has the potential to access or destroy fiber-channel connected disk storage. Two areas of concern: 1) verification that only authorized users (usually limited to only root or hpss) are granted read and write access to these resource; 2) HPSS administrators should be aware that machines, possibly owned or managed by other groups, which are added to the SAN to facilitate the use of SAN3P transfers will have access to all data on disk and tape resources. If

those systems are compromised, or there are individuals authorized for systems privileges on those particular machines, but not necessarily authorized for HPSS administrative access, there is the potential for access and/or damage to HPSS data. These are inherent limitations of SAN implementations that have not yet been addressed by the industry and cannot be remedied by HPSS.

Multiple Mover Tasks (disk only)

If ON, the Mover will allow multiple Mover tasks to access the disk device.

Disable Direct I/O (disk only)

An indication of whether this device should have Direct I/O capability disabled. The HPSS Mover specifies O_DIRECT when working with disks by default in an effort to bypass kernel caching overhead and improve performance. However, for some disk appliances, specifying O_DIRECT can disable appliance-level buffer caching and cause poor performance.

Table 7.1. Recommended settings for tape devices

Device driver	Mover executable	NO-DELAY support	Locate support	TAOS	Fast Positioning	Disable RAO
Linux Native	hpss_mvr_tcp	ON	ON	OFF	ON	OFF

PVR (tape only)

The name of the PVR that handles the removable media operations for the drive.

Drive Address

For tape drives, this is the name or address by which the PVR can access the drive. This field isn't used by the PVL or PVR for disk drives; the administrator may use this as a comment. Many sites duplicate the **Device Name** which makes it valuable as a query of drive information since it will return the device information for the disk case (for example, `lshpss -drv`).



Advice

For StorageTek robots: Drive Address configuration entries correspond to the ACS,Unit,Panel,Drive Number used by ACSLS to identify drives. For example, the first drive in a typical configuration has Drive Address 0,0,10,0.

*For SCSI PVRs: The Drive Address configuration entries correspond to their associated Drive Identifier output from the **device_scan** tool. The Drive Identifier can be mapped to the device using the **device_scan -t tape** command on the host which will control the devices.*

For IBM robotics, the full identifier string, including any interior spaces, must be present. If Identifier string is:

```
Identifier = ['T10:IBM|ULT3580-TD3      1210395617']
```

then the Drive Address is

```
T10:IBM|ULT3580-TD3      1210395617
```

*For Spectra Logic robotics, the Drive Address is the output of the **device_scan -i**. However, it must be mapped to the correct drive device as identified by **device_scan -***

***t** tape in order to configure the HPSS device path. These identifiers are close but not identical. For example, **device_scan -i** may indicate:*

```
1 = [256] ['T10:10110043|4C']
```

*while **device_scan -t** tape may indicate:*

```
Identifier = ['T10:IBM|ULTRIUM-TD8 101100434C']
```

*These two strings refer to the same device; the library has removed the product and whitespace ("IBM|ULTRIUM-TD8 ") that is displayed in the **device_scan -t** output. However, the string from **device_scan -i** is what is required, since that is what the SCSI PVR will use to detect the device. The Drive Address should be:*

```
T10:10110043|4C
```

For operator mounted drives: For manually mounted drives, the drive address string is displayed to operators when an explicit drive is selected for mounting. The drive address should therefore be a string that easily communicates to an operator the drive in question (that is, a name matching the label on the exterior of the drive).

For AML robots: Leave the Drive Address field blank.

The AML PVR is supported by special bid only.

Controller ID

An indication of the adapter/bus that is in the data path to the device.



Advice

This field is used to attempt to mount individual volumes of a set of striped media on individual controllers, when possible, to prevent contention on the data path. It is recommended that the drives that are on the same adapter/bus have the same Controller ID.

Polling Interval (tape only)

The number of seconds to wait between polling requests performed by the PVL to determine if any media is present in the drive. This field is disabled by default. Use "-1" to disable polling. Values of "0" to "14" are not valid. If the drive is an operator-mounted tape drive, then this field shouldn't be disabled.



Advice

It is recommended that this field be set to a value greater than 14 seconds for operator-mounted tape drives and disabled for library managed tape drives. In normal operations for library managed tape drives (for example, STK SL8500), drive polling should not be required. In reality there are occasions where polling may detect a non-HPSS cartridge or an unaccounted for HPSS cartridge. However, the use of the drive polling (originally in place for Operator PVRs) is associated with high overhead (one thread per drive that is polling). A less expensive alternative that addresses this condition is currently being pursued. In the meantime, if you're concerned with monitoring for unexpected cartridges in library managed tape drives,

it is recommended that you set the polling interval to a value no less than "900" (15 minutes).

Maintenance Date

The time and date when the drive last had maintenance performed on it. This date must be entered manually and is provided for administrative information purposes only.

Mounted Volume (disk only)

The eight-character name of the volume mounted on the disk drive.

Comment

This field provides a 128-character buffer in the PVL drive metadata which gives the administrator the opportunity to associate miscellaneous text with a device/drive. For example, a site may want to place a comment in this field that the drive is out of service or being used by another system.

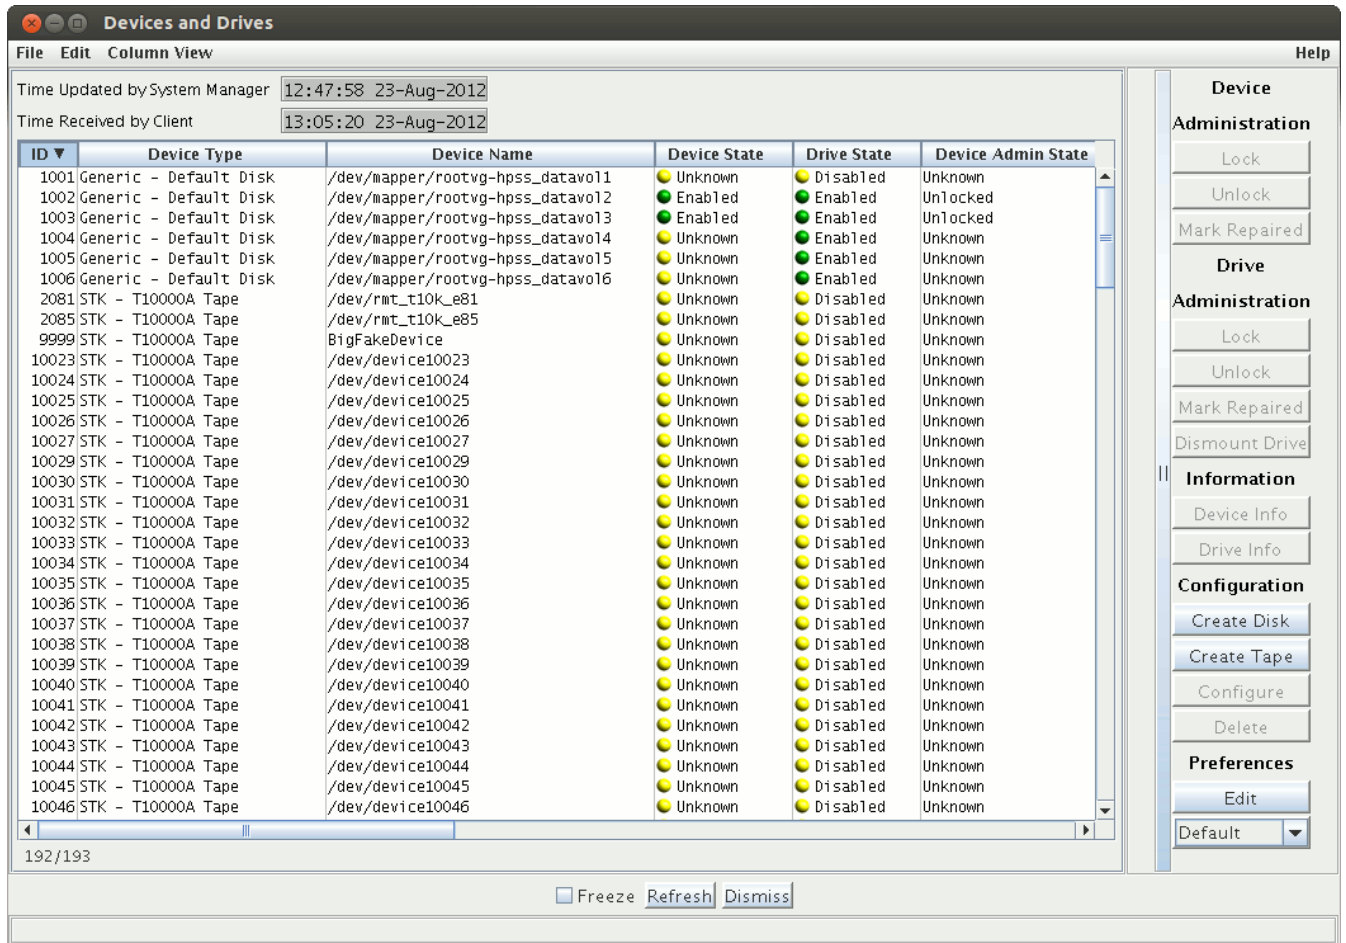
Library Unit ID (tape only)

Since library architecture varies, the method to associate related drives at a unit level will be determined through the use of a "unit id" that is placed in the drive metadata. This will allow the administrator to analyze their specific library configuration and apply their strategy as they feel fit. The use of "unit id" allows the library specifics to be abstracted from the PVL. Library details are the domain of the PVR. See **Defer Dismount Exempt Count** (PVR Server Configuration).

**Advice**

Devices configured for AWS Tape Gateway must have SCSI reservations disabled. This can be done by adding "HPSS_MVR_DISABLE_RESERVATIONS=<DeviceId>" in /var/hpss/etc/env.conf on the Mover node.

7.1.1. Devices and Drives window



This window allows you to view the list of configured Mover devices and PVL drives. It also provides a number of function buttons, which allow certain operations to be performed on devices or drives.

The *Device and Drive List Preferences* window may be used to select the device/drives which will be displayed. Select the columns to be displayed in the list from the Column View menu.

Most of the function buttons to the right of this list require that one or more device/drive entries be selected from the list. Note that some of the fields in the table may not be visible, depending on the Column View menu settings.

Field descriptions.

Device and Drive List

ID

The unique device/drive ID number defined when the device/drive is configured.

Device Type

The type of the device.

Device Name

The name used by the Mover to access the device, usually the pathname of a UNIX device special file (such as `/dev/rmt0`).

Device State

The current operational state of the device, as reported by its controlling Mover. The possible states are:

- `Enabled` indicates the device is working normally.
- `Suspect` indicates the Mover has detected errors on the device, but it is still functioning.
- `Disabled` indicates the device is locked, which makes it unavailable for use by HPSS.
- `Unknown` indicates the state of the device is not known to SSM; this is usually caused by the controlling Mover being down or disconnected from SSM.

Device Admin State

The current administrative state of the device, as reported by its controlling Mover. The possible states are:

- `Locked` indicates the device is locked and unavailable for use by HPSS.
- `Unlocked` indicates the device is unlocked and available for use by HPSS.
- `Unknown` indicates the state of the device is not known to SSM; this is usually caused by the controlling Mover being down or disconnected from SSM.

Mover

The descriptive name of the Mover that controls the device.

Mover Host

The hostname of the Mover node (the node where the Mover runs).

Drive Address

The name used by the PVL/PVR to access the drive.

Drive Flag

This field describes possible actions that the PVL needs to perform for the drive. Possible flags are:

- `Clear` indicates nothing needs to be done.
- `Modified` is set by SSM when the drive has been modified. This value is reserved for use in a planned feature. It is not used at this time.
- `PVR/Mover Notify Pending` indicates the PVL needs to notify the associated PVR and Mover that the drive has been created or deleted.
- `PVR Notify Pending` indicates the PVL needs to notify the associated PVR that the drive has been created or deleted.
- `Mover Notify Pending` indicates the PVL needs to notify the associated Mover that the drive has been created or deleted.
- `Abort PVR/Mover Notify` indicates the PVL is aborting a pending notification.

Drive State

The current operational state of the drive, as reported by the PVL. The possible states are:

- **Enabled** indicates the drive is working normally.
- **Disabled** indicates the drive is not available for use by HPSS. A drive may be disabled if it is locked by an SSM user, if an error occurred when mounting or dismounting a cartridge on the drive, or if the Core Server was unable to connect to the Mover after a tape was mounted on the drive.
- **Broken** indicates the PVL itself has detected a fatal error and is shutting itself down.
- **Unknown** indicates the state of the device is not known to SSM; this is usually caused by the PVL being down or disconnected from SSM.

Drive Admin State

The current administrative state of the drive, as reported by the PVL. The possible states are:

- **Locked**, indicating the drive is locked and unavailable for use by HPSS.
- **Unlocked**, indicating the drive is unlocked and available for use by HPSS.
- **Unknown**, indicating the state of the drive is not known to SSM; this is usually caused by the PVL being down or disconnected from SSM.

Library Unit ID (tape only)

Since library architecture varies, the method to associate related drives at a unit level will be determined through the use of a "unit id" that is placed in the drive metadata. This will allow the administrator to analyze their specific library configuration and apply their strategy as they feel fit. The use of "unit id" allows the library specifics to be abstracted from the PVL. Library details are the domain of the PVR. See **Defer Dismount Exempt Count** (PVR Server Configuration).

Comment

This field provides 128 characters of storage in the PVL drive metadata in which the administrator may store text messages associated with a device/drive. For example, a site may want to place a comment in this field that the drive is out of service or being used by another system.

PVR

The descriptive name of the PVR that handles removable media operations for the drive. This will be blank if the device/drive does not support removable media.

Mounted Volume

The name of the volume that is currently mounted as reported by the PVL.

Bytes Read

Number of bytes read from the device/drive as reported by the Mover.

Bytes Written

Number of bytes written to the device/drive as reported by the Mover.

State

The overall operational state of the device/drive. If this is not **Enabled**, check the individual **Device State** and **Drive State** fields to help determine why.

Device Administration buttons.

This group of buttons affects selected Mover devices. All the buttons are disabled unless one or more devices are selected (see figure above).

Lock

Locking the selected devices serves no functional purpose. When the device is locked, the information is not transferred to the PVL and thus the PVL will continue to schedule the associated PVL drive (see **Lock** button in the **Drive Administration Buttons** below). It is also not transferred to the Core Server which will continue to read, write, create, or mount the volume (see **VV Condition** in the *Core Server Volume Disk/Tape* window). This leads to many red Alarms. After clicking the Device Administration **Lock** button, you are prompted to confirm your request before it is actually executed. The window's message line reports when locking begins. When the request is complete, another message tells how many requests were successful.

Unlock

Unlock the selected devices, making them available to the Mover. You are prompted to confirm your request before it is actually executed. The window's message line reports when unlocking begins. When the request is complete, another message tells how many requests were successful.

Mark Repaired

Marks the selected devices as "repaired". Sometimes the device states (such as operational state) will continue to indicate an error condition after the cause of the error has been fixed. If so, you can mark a device repaired to instruct the Mover to clear its error states. Note that this does nothing, either in hardware or software, to actually repair an error condition. Also, if you mark a device repaired when it still has a problem, the error states will be cleared but may quickly return to their previous values.

You are prompted to confirm your request before it is actually executed. The window's status bar displays a message when marking begins, and when the request is complete, another message tells how many requests were successful.

Drive Administration buttons.

This group of buttons affects selected PVL drives. All the buttons are disabled unless one or more drives are selected (see figure above).

Lock

Lock the selected drives, making them unavailable to HPSS. When a drive is locked, the PVL will no longer schedule the PVL drive. When locking a tape drive due to a cartridge problem (such as for a stuck tape), we recommend canceling the PVL Job associated with the cartridge. We also recommend changing the Condition of the tape volume in the Core Server that owns it so that no attempts will be made to perform I/O operations on the volume. The "Down" condition is provided for this purpose. Locking the PVL drive does not automatically change the volume's Condition. After clicking the Drive Administration **Lock** button, you are prompted to confirm your request before it is performed. The window's message line reports when locking begins. When the request is complete, another message tells how many requests were successful.

Unlock

Unlock the selected drives, making them available to HPSS. You are prompted to confirm your request before it is actually performed. The window's message line reports when unlocking begins. When the request is complete, another message tells how many requests were successful.

Mark Repaired

Notify the PVL and Mover to clear any error states for the selected drives. Sometimes the drive states (such as operational state) will continue to indicate an error condition after the cause of the error has been corrected. If so, you can mark a drive repaired to instruct the Mover or PVL to clear its error states. Note that this does nothing, either in hardware or software, to actually repair an error condition. Also, if you mark a drive repaired when it still has a problem, the error states will be cleared but may quickly return to their previous values.

You are prompted to confirm your request before it is actually executed. The window's status bar displays a message when marking begins, and when the request is complete, another message tells how many requests were successful.

Dismount

Dismount the selected drives. You are prompted to confirm your request before it is actually executed. The window's status bar displays a message when dismounting begins. When the request is complete, another status message tells how many requests were successful.

Information buttons.

This group of buttons retrieves the device-managed or drive-managed object for one or more selected list items. All the buttons are disabled unless one or more devices/drives are selected.

Device Info

Open the *Mover Device Information* window for the selected item. If the selected item's Mover cannot be contacted, this button will be desensitized. While this window is open, it will continue to update reflecting the latest device data from the Mover.

Drive Info

Open the *PVL Drive Information* window for the selected item. If the PVL cannot be contacted, this button will be desensitized. While this window is open, it will continue to update reflecting the latest drive data from the PVL.

Configuration buttons.

This group of buttons allows you to do device/drive configuration.

Create Disk

This button is always active. Clicking on it opens the *Disk Device Configuration* window in Add mode, allowing you to create a new disk device and drive.

Create Tape

This button is always active. Clicking on it opens the *Tape Device Configuration* window in Add mode, allowing you to create a new tape device and drive.

Configure

This button is active if one or more devices/drives are selected. Clicking the button opens the *Disk/Tape Device Configuration* window for the selected devices and drives, allowing you to view, delete, clone, or update the configurations.

Delete

This button is active if one or more device/drives is selected from the list. A dialog will be displayed for each device/drive selected asking the user to confirm that the device/drive configuration is to be deleted.

Preferences.

See *Section 7.1.1, “Devices and Drives window”* for information on changing preferences for the *Devices and Drives* window.

7.1.2. Enable variable block sizes for tape devices

All tape devices that will be used for HPSS data must be set to handle variable block sizes (to allow for the ANSI standard 80-byte volume label and file section headers). This requirement is described in detail in the *Tape devices* subsection of the *HPSS Installation Guide* (within the *Mover* section of the *HPSS planning* chapter). Refer to that section to be certain your tape devices are configured correctly.

7.1.3. Changing a drive's configuration

Devices and drives can be updated dynamically (that is, while the PVL, associated Movers, and associated PVRs (for tape) are up). Many of the same required conditions for drive deletion must be met. The drive must be locked without any activity (drive state = FREE). Additionally for disk, the volume must be unallocated. However, there are some drive attributes for which the configuration metadata can be updated dynamically via the *PVL Drive Information* window and there are some device attributes for which the configuration metadata is not updated dynamically but the server "in memory" values can be changed dynamically (but not preserved when the Mover is restarted) via the *Mover Device Information* window without the drive locking/free state step.

The device and drive configuration entry can be updated using the *Disk Device Configuration* or *Tape Device Configuration* window. The *Disk Device Configuration* and *Tape Device Configuration* windows are updated dynamically. These windows are opened by clicking the **Configure** button on the *Devices and Drives* window. Changing any settable field on the disk/tape device configuration windows requires all servers involved in the update to be running:

- the Mover that is currently managing the device
- the new Mover that will manage the device (if changing)
- the PVR that is currently managing cartridges for this drive (if tape)
- the PVR will manage cartridges for this drive (if tape and if changing)

An option exists which allows the PVR to make a configuration update without the associated servers running. This is intended for an abnormal condition and should be applied sparingly and with caution. If the PVL cannot communicate with a Mover, the PVL communicates the condition to the SSM. The SSM will prompt the administrator with an override option which if set will direct the PVL to update metadata without successfully communicating to the associated Mover. It's expected that this option may be used in a failover scenario where a Mover platform has crashed and drives are moved to another platform while the failed platform is repaired or replaced. The HPSS drive move would be preceded by drive hardware changes to make the drives available at the moved platform.

Some PVL drive configuration attributes can be updated dynamically using the *PVL Drive Information* window (see *Section 7.2.2, “PVL Drive Information window”*). The settable fields in this window are updated dynamically (that is, saved to metadata and used by the PVL upon successful update).

It is also possible to make temporary updates to some Mover device attributes parameters by changing the settable fields in the *Mover Device Information* window (see *Section 7.2.1, “Mover Device Information window”*) instead of the configuration window. These changes stay in effect until the Mover is recycled.

7.1.4. Deleting a drive's configuration

Devices and drives are deleted dynamically. This means that they must be deleted while the system is running. Specifically the PVL must be up; additionally it is expected that the associated Mover and associated PVR (for tape) are also up since the PVL needs to notify these servers about the deletion. If the PVL is unable to successfully notify the associated Mover and/or PVR about the deleted device/drive, it will retry several times; each failure will be logged as an ALARM. After the last failed attempt to notify the associated Mover and/or PVR, the PVL will abandon retrying; the administrator is advised to bounce the abandoned associated Mover (if it was up during the failed retries) and/or to reinitialize the abandoned associated PVR (if it was up during the failed retries) to ensure that the device/drive is completely deleted and thus can no longer be used.

Likewise, a previously added device/drive cannot be deleted until the PVL has successfully notified the associated Mover and PVR (for tape) about the previous add. If the administrator attempts to delete a device/drive with a previously added Drive ID that is still pending add notification, the administrator will get a BUSY failure and the pending notification will be aborted. An ALARM will be logged informing the administrator that the delete failed since the Mover/PVR add/create notification is pending. Another ALARM will be logged informing the administrator that the previous drive create notification will be aborted and to try later; wait a minute or so before retrying in order to give the PVL time to abort the notification. The **Drive Flag** field in the *Devices and Drives* window and *PVL Drive Information* window will report whether or not a particular drive is pending Mover/PVR notification; of course, this is only viewable while the drive exists.

There are a number of situations in which the PVL won't allow the device/drive to be deleted:

- If the device/drive is still attempting to notify the Mover/PVR about being added
- If the device/drive is in the process of aborting Mover/PVR notification
- If the drive is in use by the PVL
- If the drive is not locked
- For disk: If storage resources haven't been deleted from the disk device/drive
- For disk: If the physical volume hasn't been exported from the disk device/drive

If it is only necessary to take a drive out of service for a finite period of time, the administrator should follow the instructions in *Section 7.1.1, “Devices and Drives window”* to lock the drive and prevent its use rather than permanently removing the drive as described in this section.

The Device and Drive configuration entry can be deleted using the *Devices and Drives* list window (see *Section 7.1.1, “Devices and Drives window”*). Select the correct device/drive configuration and click on the **Delete** button to delete the configuration. Once the drive configuration is deleted, HPSS will no longer be able to use the drive's resources.

7.2. Monitoring devices and drives

After the devices/drives are configured and are made available for service, the device/drive status may be monitored via several SSM windows.

When a device/drive experiences a problem, the PVL and/or the Mover issues one or more alarms to inform SSM users of the problem and to notify SSM of the changes in the device/drive state. SSM reports the state change on the **Devices and Drives** status field on the *HPSS Health and Status* window. SSM also reports the new state on the *Devices and Drives* list window, the *Mover Device Information* window, and the *PVL Drive Information* window if they are being displayed at the time.

If an HPSS device or drive needs to be locked or unlocked or if any device/drive configuration needs to be done, the *Devices and Drives* list window is the first stop (for detailed information, see *Section 7.1.1, “Devices and Drives window”*). It facilitates drive locking and unlocking, and it also lets administrators access the following windows:

- *Disk/Tape Device Configuration* is used to configure, modify, or delete device and drive configurations. See for more information.
- *Mover Device Information* is used to view or update the state of the Mover device. See *Section 7.2.1, “Mover Device Information window”* for more information.
- *PVL Drive Information* is used to view or update the state or configuration of the PVL drive. See *Section 7.2.2, “PVL Drive Information window”* for more information.

7.2.1. Mover Device Information window

Mover Device Information

File Edit Help

Basic Information

Device Name

/dev/tape/by-id/scsi-3500507630f316809

Device ID

23

Device Type

IBM - 3592 E07/EH7 (TS1140) Tape

Mover

Mover (camille)

Administrative State

Unlocked

Operational State

Suspect

Usage State

Idle

Device Flags

☒ Read Enabled

☒ NO-DELAY Support

☒ Removable Media Support

☒ Write Enabled

☐ TAOS Support

☒ Locate Support

☒ Fast Positioning

☐ Disable RAO

State

Set State

Volume Flags

0

Number Of Errors

1

Reset

Block Size

0

TaosFD

0

Volume ID

LBPFlags

0

Bytes Read

2,148,532,224

Reset

Bytes Written

0

Reset

☐ Freeze

Update

Refresh

Dismiss

Retrieving data - Succeeded

Mover Device Information
_ □ ×

File Edit
Help

Basic Information

Device Name	<input type="text" value="/opt/disk/disk1"/>
Device ID	<input type="text" value="5"/>
Device Type	Generic - Default Disk
Media Block Size	<input type="text" value="4096"/>
Capacity	<input type="text" value="1TB"/>
Starting Offset	<input type="text" value="0"/>
Mover	<input type="text" value="Mover (camille)"/>
Administrative State	<input type="text" value="Unlocked"/> ▼
Operational State	● Enabled
Usage State	● Idle

Device Flags

☒ Read Enabled

☒ Write Enabled

☐ SAN3P-transfer Enabled

☒ Multiple Mover Tasks

☒ Disable Direct I/O

State	<input type="text"/>
Set State	<input type="text"/>
Volume Flags	<input type="text" value="0"/>
Number Of Errors	<input type="text" value="0"/> Reset
Block Size	<input type="text" value="4096"/>
TaosFD	<input type="text" value="0"/>
Volume ID	<input type="text"/>
LBPFlags	<input type="text" value="0"/>
Bytes Read	<input type="text" value="0"/> Reset
Bytes Written	<input type="text" value="0"/> Reset

Retrieving data - Succeeded

The *Mover Device Information* window reports the current statistics for the device, such as the workload history of the device since the startup of the controlling Mover. The *Mover Device Information* window can also be used to lock and unlock a Mover device. Additionally, it can be used to control the I/O aspects of the device.

Note: locking the Mover device generally is not helpful; see *Section 7.1.1, “Devices and Drives window”*.

Changes to the **Administrative State** cause a confirmation window to pop up. If the user confirms the request, modifications made to all fields on the window will be sent to the Mover and become effective immediately.

Field descriptions.

Device Name

The name by which the Mover accesses the device; usually the pathname of a UNIX device special file (such as `/dev/rmt0`).

Device ID

The ID number unique to this device. The same ID number also refers to the corresponding PVL drive.

Device Type

The HPSS media type which describes the physical device.

Media Block Size (disk only)

The block size for the device.

Capacity (disk only)

The size of the disk device in bytes.

Starting Offset (SAN3P disk only)

The offset in bytes from the beginning of the disk logical volume at which the Mover will begin using the volume. The space preceding the offset will not be used by HPSS.

Mover

The descriptive name of the Mover which controls the device.

Administrative State

This field allows you to modify the state of the device. Click on the option menu button to pop up a list of valid states. These are:

- `Locked` makes the device unavailable for HPSS requests.
- `Unlocked` makes the device available for HPSS requests.
- `Mark Repaired` tells the Mover to clear any error status for the device. This can be useful if you think a problem has been fixed, but the Mover is unaware of it. This does not do anything, in hardware or software, to fix a problem; it only clears error indicators.

Operational State

Indicates the device's ability to handle HPSS requests. Possible values for this field are:

- `Enabled`, indicating the device is available for HPSS requests.
- `Disabled`, indicating the device is unavailable for HPSS requests, possibly caused by setting the "Locked" Administrative State.

Usage State

Indicates the state of the Mover's control over the device. Possible values for this field are:

- `Active` indicates a Mover task currently has control of the device.
- `Idle` indicates no Mover task is controlling the device.

Device Flags

This group of buttons defines various device characteristics. When a button is ON, it means that the corresponding flag is set.

Read Enabled

An indication of whether the device is available for reading.

Write Enabled

An indication of whether the device is available for writing.

Locate Support (tape only)

An indication of whether the device supports a high speed (absolute) positioning operation.



Advice

This option is supported for IBM 3590, 3590E, 3580, 3592, StorageTek 9840, 9940, T10000 and GY-8240 devices.

NO-DELAY Support (tape only)

An indication of whether the device supports opening the device with no delay flag set, while allowing tape I/O operation after the open.



Advice

On some tape devices, this will allow for a quicker polling operation when no tape is presently loaded in the device. This field is meaningful for tape devices only.

TAOS Support (tape only)

An indication of whether the device is set up to use the Spectra Logic TAOS feature. TAOS enables an RAO-like interface for devices which do not support the RAO command set.



Advice

What is TAOS? [https://spectralogic.com/wp-content/uploads/white_paper_increasing_file_read_throughput_on_tape_with_taos.pdf]

TAOS Support should only be enabled for devices which have already been set up for TAOS with a Spectra Logic library. Discuss your data use cases with Spectra Logic and HPSS before enabling, as TAOS can result in a lower performance in certain scenarios. The TAOS device setup can be confirmed by using `device_scan` or `lsscsi` on the Mover device. Verify that Spectra Logic supports TAOS on the device type prior to enabling this flag.

Fast Positioning (tape only)

An indication of whether the device should skip section header verification when positioning via LBA.

Disable RAO (tape only)

An indication of whether this device should have RAO (Recommended Access Order) capability disabled. This field has no effect for devices which do not support RAO.

Removable Media Support (tape only)

An indication of whether the device supports removable media.

SAN3P-transfer Enabled (disk only)

If ON, SAN3P data transfers to this device will be supported.



Warning - *There is a security vulnerability associated with the use of SAN3P. If a user is root on a machine which has access to the SAN (for example, a client machine) then that user has the potential to access or destroy fiber-channel connected disk storage. Two areas of concern: 1) verification that only authorized users (usually limited to only root or hpss) are granted read and write access to these resource; 2) HPSS administrators should be aware that machines, possibly owned or managed by other groups, which are added to the SAN to facilitate the use of SAN3P transfers will have access to all data on disk and tape resources. If those systems are compromised, or there are individuals authorized for systems privileges on those particular machines, but not necessarily authorized for HPSS administrative access, there is the potential for access and/or damage to HPSS data. These are inherent limitations of SAN implementations that have not yet been addressed by the industry and cannot be remedied by HPSS.*

Multiple Mover Tasks (disk only)

If ON, the Mover will allow multiple Mover tasks to access the disk device.

Disable Direct I/O (disk only)

An indication of whether this device should have Direct I/O capability disabled. The HPSS Mover specifies `O_DIRECT` when working with disks by default in an effort to bypass kernel caching overhead and improve performance. However, for some disk appliances, specifying `O_DIRECT` can disable appliance-level buffer caching and cause poor performance.

State

A set of flags describing the current state of the device. The flag states are presented as a list of flag descriptions for each flag which is currently ON:

- `Mounted` indicates a media volume is mounted in the device.
- `Open for Read` indicates the device is open for reading.
- `Open for Write` indicates the device is open for writing.
- `Busy` indicates the device is in use.
- `Error Encountered` indicates an error was encountered on the device.
- `Hit End-Of-Tape` indicates the end of a tape was detected.
- `Client Off Midblock` indicates the device position is in the middle of a block.
- `Data to be Synced` indicates there are data that need to be flushed.

- Last Op Read indicates the last operation on the drive was a read.
- Logical Block Support indicates the device supports Logical Block Addressing.

Set State

A set of flags showing device state changes which are pending. The flag states are presented as a list of flag descriptions for each flag which is currently ON:

- Toggle Read Mode - Toggle the current read mode value on the device.
- Toggle Write Mode - Toggle the current write mode value on the device.
- Toggle In-Use Mode - Release the device from the in use state.
- Reset Error Count - Reset the device **Number of Errors** field
- Reset Bytes Read - Reset the **Bytes Read** field.
- Reset Bytes Written - Reset the **Bytes Written** field.
- Set to Locked - Set the device **Administrative State** to locked.
- Set to Unlocked - Set the device **Administrative State** to unlocked.
- Set to Repaired - Clear any error status for the device.

Volume Flags

Flags which further describe that describe the volume type.

Number of Errors

The number of errors that have occurred on this device since the last time this field was cleared. Click on the **Reset** button to the right of the field to reset the count to zero.

Block Size

The size, in bytes, of each data block on this device.

Volume ID

The label (name) of the volume which is currently mounted on this device. If no volume is currently mounted, this field will be blank.

Bytes Read

The number of bytes that have been read from this device since the last time this field was cleared. Click on the **Reset** button to the right of the field to reset the count to zero.

Bytes Written

The number of bytes that have been written to this device since the last time this field was cleared. Click on the **Reset** button to the right of the field to reset the count to zero.

Operational notes

- Setting device attributes through this window will not cause the configuration metadata to be updated.

- The device statistics fields (number of bytes read, number of bytes written, number of errors) are all initialized to zero when the Mover is initialized.
- When modifying the **Administrative State**, toggling the device flags or resetting the device statistics, changes do not take place until the device is no longer in use. These modifications will be pending until the device is released and no longer busy.

7.2.2. PVL Drive Information window

The screenshot shows the 'PVL Drive Information' window. It has a menu bar with 'File', 'Edit', and 'Help'. The main area contains the following fields and controls:

Drive ID	8
Drive Type	Generic - Default Disk
Mover	Mover (malibu)
Administrative State	Unlocked
Operational State	Enabled
Usage State	Busy
Drive State	In Use
Drive Flag	Clear
Controller ID	8
Mounted Volume	MA000100
Maintenance Date	(ex. Jan 31, 2009 10:59:30 AM)
Drive Error Count	0 <input type="button" value="Reset"/>
Mounts Since Last Maintenance	26 <input type="button" value="Reset"/>
Comment	

At the bottom, there is a 'Freeze' checkbox and three buttons: 'Update', 'Refresh', and 'Dismiss'. A status bar at the very bottom displays 'Retrieving data - Succeeded'.

The screenshot shows a window titled "PVL Drive Information" with a menu bar containing "File", "Edit", and "Help". The window contains the following fields and controls:

- Drive ID:** Text field with value "4".
- Drive Address:** Text field with value "/dev/drive4".
- Drive Type:** Text field with value "Generic - Default Tape".
- Mover:** Text field with value "Mover (hpss-dev02)".
- PVR:** Text field with value "Operator PVR".
- Administrative State:** Dropdown menu with "Locked" selected.
- Operational State:** Radio button labeled "Disabled" (selected).
- Usage State:** Radio button labeled "Idle" (selected).
- Drive State:** Text field with value "Free".
- Drive Flag:** Text field with value "Clear".
- Controller ID:** Text field with value "4".
- Polling Interval:** Text field with value "15" and label "seconds (-1 to disable)".
- Mounted Volume:** Text field (empty).
- Maintenance Date:** Text field (empty) with example text "(ex. Dec 31, 2013 10:59:30 AM)".
- Drive Error Count:** Text field with value "0" and a "Reset" button.
- Mounts Since Last Maintenance:** Text field with value "0" and a "Reset" button.
- Comment:** Text area (empty).
- Library Unit ID:** Text field with value "500".

At the bottom of the window, there are four buttons: "Freeze", "Update", "Refresh", and "Dismiss". A status bar at the very bottom displays the message "Retrieving data - Succeeded".

This window allows you to view and update the information associated with an HPSS drive. The *PVL Drive Information* window is typically used to lock and unlock drives since newly configured drives are locked by default and must be unlocked to be used. It may also be used to determine which volume is mounted on the drive when the drive reports a mount error condition.

Any changes made to fields on this window are saved in metadata and sent directly to the PVL and thus are effective immediately. Thus the changes are dynamic.

Field descriptions.

Drive ID

The ID number unique to this drive. The same ID number also refers to the corresponding Mover device.

Drive Address (tape only)

The name/address by which the drive's PVR can access it. This information is used in communications between the PVL and the PVR that controls removable media operations for the drive.

Drive Type

The HPSS drive type which describes the physical drive.

Mover

The descriptive name of the Mover which moves data to and from this drive.

PVR (tape only)

The descriptive name of the PVR used to control this drive. This field is only meaningful for tape drives.

Administrative State

This field allows you to modify the state of the drive. The options are:

- `Locked` makes the drive unavailable for HPSS requests.
- `Unlocked` makes the drive available for HPSS requests.
- `Mark Repaired` tells the PVL to clear any error status for the drive. This can be useful if you think a problem has been fixed, but the PVL is unaware of it. This does not do anything, in hardware or software, to fix a problem; it only clears error indicators.

Operational State

Indicates the drive's ability to handle HPSS requests. Possible values for this field are:

- `Enabled` indicates the drive is available for HPSS requests.
- `Disabled` indicates the drive is unavailable for HPSS requests; possibly caused by setting the `Locked` administrative state (see above) or if an error occurred when mounting or dismounting a cartridge on the drive.
- `Broken` will appear only when the PVL is in the midst of shutting itself down because of a fatal error.
- `Unknown` indicates the state of the device is not known to SSM; this is usually caused by the PVL being down or disconnected from SSM.

Usage State

Indicates the state of the PVL's control over the drive. Possible values for this field are:

- `Active` indicates the drive is in use.
- `Idle` indicates the drive is not in use.
- `Busy` indicates the drive is busy.
- `Unknown` indicates the state is not known to SSM.

Drive State

The current state of the drive as reported by the PVL. Possible values for this field:

- `Clear` indicates the drive has been created, but the associated Mover and PVR (if tape) haven't successfully been notified.
- `In Use` indicates a volume is currently mounted or in the process of being mounted on the drive.
- `Free` indicates the drive is unoccupied.

- `Dismount Pending` indicates the drive is in the process of dismounting a volume.
- `Deleted` indicates the drive has been deleted (metadata no longer exists).

Drive Flag

This field describes possible actions that the PVL needs to perform for the drive. Possible flags are:

- `Clear` indicates nothing needs to be done.
- `Modified` is set by SSM when the drive has been modified. It was added for a future feature; currently unsupported.
- `PVR/Mover Notify Pending` indicates the PVL needs to notify the associated PVR and Mover that the drive has been created or deleted.
- `PVR Notify Pending` indicates the PVL needs to notify the associated PVR that the drive has been created or deleted.
- `Mover Notify Pending` indicates the PVL needs to notify the associated Mover that the drive has been created or deleted.
- `Abort PVR/Mover Notify` indicates the PVL is aborting a pending notification.

Controller ID

An indication of the adapter/bus that is in the data path to the device. This field is used to attempt to mount individual volumes of a set of striped media on different controllers, when possible, to prevent contention on the data path.

Polling Interval (tape only)

The number of seconds to wait between polling requests performed by the PVL to determine if any media is present in the drive. This field is disabled by default. Use "-1" to disable polling. Values of "0" to "14" are not valid. If the drive is an operator-mounted tape drive, then this field shouldn't be disabled.



Advice

It is recommended that this field be set to a value greater than 14 seconds for operator-mounted tape drives and disabled for library managed tape drives. In normal operations for library managed tape drives (for example, STK SL8500), drive polling should not be required. In reality there are occasions where polling may detect a non-HPSS cartridge or an unaccounted for HPSS cartridge. However, the use of the drive polling (originally in place for Operator PVRs) is associated with high overhead (one thread per drive that is polling). A less expensive alternative that addresses this condition is currently being pursued. In the meantime, if you're concerned with monitoring for unexpected cartridges in library managed tape drives, it is recommended that you set the polling interval to a value no less than "900" (15 minutes).

Mounted Volume

The volume that is currently mounted on this drive.

Maintenance Date

The time and date when the drive last had maintenance performed on it. This date must be entered manually and is provided for administrative information purposes only.

Drive Error Count

This field records the number of consecutive Drive Errors (set via the **Retry Mount Time Limit** mechanism) which have occurred on the drive. When the number of errors equal or exceed the **PVR Drive Error Limit**, the drive will automatically be locked by the PVL. The field can be reset to zero administratively using the **Reset** button. A successful mount or PVL recycle will also zero the field.

Mounts Since Last Maintenance

The number of times a volume has been mounted on this drive since the last maintenance. Clicking on the **Reset** button to the right of this field will reset the value to zero.

Comment

This field provides a 128-character buffer in the PVL drive metadata which gives the administrator the opportunity to associate miscellaneous text with a device/drive. For example, a site may want to place a comment in this field that the drive is out of service or being used by another system.

Library Unit ID (tape only)

Since library architecture varies, the method to associate related drives at a unit level will be determined through the use of a "unit id" that is placed in the drive metadata. This will allow the administrator to analyze their specific library configuration and apply their strategy as they feel fit. The use of "unit id" allows the library specifics to be abstracted from the PVL. Library details are the domain of the PVR. See **Defer Dismount Exempt Count** (PVR Server Configuration).

Related information.

Section 17.1.2, "Tape management".

7.3. Changing device and drive state

The administrative state of a device or drive can be set to Unlocked or Locked. This controls whether HPSS can access the drive. Changing the state of a device or drive can be accomplished via the *Devices and Drives* list window.

Notice that there are two sets of **Lock**, **Unlock** and **Mark Repaired** button groups on the *Devices and Drives* window. The first group is titled **Device Administration** and the second group is titled **Drive Administration**. The first group is working at the Mover device level and the second group is working at the PVL drive level.

Locking *devices* serves no functional purpose. When the device is locked, the information is not transferred to the PVL and thus the PVL will continue to schedule the associated PVL drive. It is also not transferred to the Core Server which will continue to read and write the volume (see **VV Condition** in the *Core Server Volume Disk/Tape* window) that is mounted in the drive.

When a *drive* is locked, the PVL will no longer schedule the PVL drive. When locking a tape drive due to a cartridge problem (such as for a stuck tape), it is beneficial to also cancel the PVL job associated with the cartridge. Additionally, locking the drive is not transferred to the Core Server,

which will continue to read and write the volume (see **VV Condition** in the *Core Server Volume Disk/Tape* window) unless you change the **VV Condition** to DOWN (or EOM) which is advised while diagnosing the stuck tape.

7.3.1. Unlocking a drive

The **Administrative State** of a drive must be "unlocked" to allow HPSS to make use of a device/drive. Before unlocking a drive, ensure that its hardware is functional and is fully configured into its host's operating system and into HPSS. Configuration of an HPSS drive is described in *Section 7.1, "Configure a new device and drive"*.

From the *Devices and Drives* window (see *Section 7.1.1, "Devices and Drives window"*), select the desired device/drive entries and then click on the **Unlock** button from the Drive button group. The drive can also be unlocked by bringing up the *PVL Drive Information* window and setting its **Administrative State** to Unlocked.

7.3.2. Locking a drive

To lock a drive, go to the *Devices and Drives* window (see *Section 7.1.1, "Devices and Drives window"*), select the desired device/drive entries and then click on the **Lock** button from the Drive button group. The drive can also be locked by bringing up the *PVL Drive Information* window and setting its **Administrative State** to Locked.

Locking a drive disallows its use by HPSS. Changing a drive state to locked will ensure that the drive will not be used for new mounts.

Always lock the PVL drive instead of the Mover device. Mover devices do not usually need to be locked or unlocked, and the capability to lock and unlock them is provided for troubleshooting purposes only.

Locking a tape drive will not affect an active job which has a volume mounted. Once that job has completed, the drive will be dismounted and no further mounts will take place. This may be useful when preventative maintenance is required for an operating drive.

Locking a disk drive has little effect since disks are logically mounted when the PVL initializes and are not usually dismounted; however, a disk drive must be in the locked state to be deleted.

7.3.3. Repairing the state of a device or drive

A drive can enter an error or suspect state as reported by the PVL, Mover, or both. After a drive has entered one of these abnormal states, it can be repaired to return it to a normal state.

From the *Devices and Drives* window (see *Section 7.1.1, "Devices and Drives window"*), select the desired device/drive entries and then click on the **Mark Repaired** button appropriate for that device/drive type. Another way to change a device/drive state to "repaired" is to bring up the *PVL Drive Information* window or the *Mover Device Information* window and changing the Administrative State to Mark Repaired.

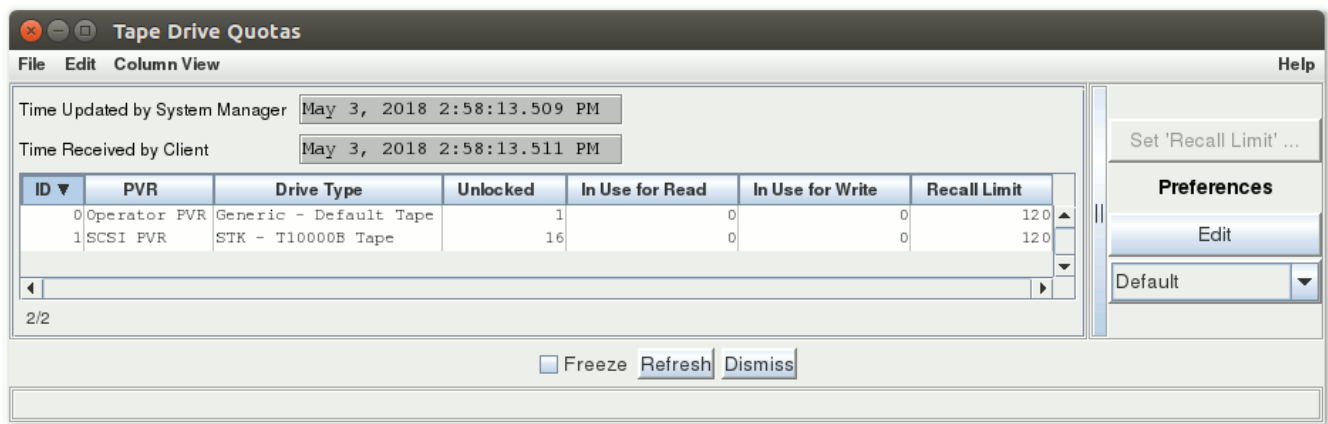
Repairing the state of a device/drive is only an instruction to the server to reset the displayed device/drive state value. It does not correct any underlying problems that might still exist. Rather, it is

a means by which the administrator can notify the server that the underlying problem has been addressed outside of the server. It is used for problems that the server cannot fix by itself such as a drive being offline or a tape getting stuck in a drive.

7.3.4. Resetting drive statistics

Both the PVL and Mover maintain statistics related to a drive. The PVL maintains an error count as well as a count of the number of times a cartridge is mounted. The Mover maintains the count of the number of errors encountered and the number of bytes read and written on the drive. These values can all be reset to zero by the **Reset** button located next to the statistic field on the appropriate window.

7.3.5. Tape Drive Quota list window



This window allows you to manage the tape drive recall limit by PVR and tape drive type.

Field descriptions.

ID

The ID associated with this device/drive.

PVR

The descriptive name of the PVR.

Drive Type

The tape drive type.

Unlocked

The number of tape drives of this type in this PVR that are currently unlocked.

In Use for Read

The number of tape drives of this type in this PVR that are currently in use for read activity.

In Use for Write

The number of tape drives of this type in this PVR that are currently in use for write activity.

Recall Limit

The maximum number of drives of this type in this PVR that can be used at the same time for data recall. A value of -1 disables this feature.

Chapter 8. Volume and storage management

This chapter describes the procedures for adding, removing, monitoring, and managing storage space in the HPSS system.

The basic unit of storage which can be added to the HPSS system is the volume. Before volumes can be added to HPSS, the underlying configuration structures must be created to support them. These structures include:

- Storage classes, hierarchies, classes of service, migration policies, purge policies, and, optionally, file families. See the chapter *Chapter 6, Storage configuration* for the procedures for creating these structures.
- Device configurations. See the chapter *Chapter 7, Device and drive management* for the procedures for creating these structures.

8.1. Adding storage space

Adding storage space to the HPSS system means creating new virtual volumes and adding them to a storage class. Creating new virtual volumes in a storage class is a two-step process:

1. **Import volumes.** In the first step, the physical volumes that will be gathered into the new Core Server virtual volumes are imported by the PVL.
2. **Create resources.** In the second step, metadata for the new volumes is created by the Core Server; this process is called creating storage resources.

8.1.1. Importing volumes into HPSS

Importing a physical volume (PV) into HPSS is the process of adding a tape cartridge or disk volume to the PVL and labeling it with an HPSS volume label. Importing also makes the volume known to the PVR and PVL, and creates the appropriate metadata records. Cleaning cartridges are managed by the tape library directly and should not be imported into the HPSS system.



To import a tape cartridge into HPSS, the administrator must be familiar with the operation of the PVR to which the cartridge will be added because the physical cartridge injection process differs among PVRs.

- **Operator** - No cartridge injection step is necessary (or possible).
- **AML** - When importing new cartridges into HPSS, the cartridges must be entered into the AML Insert/Eject/Foreign (I/E/F) port before any HPSS import operations are performed. The HPSS AML PVR moves the cartridges into the towers or shelves before importing into HPSS. Users can then use the HPSS *Import Tape Volumes* window to import the tapes into HPSS; users do not need to issue the **dasadmin insert** command prior to importing. See the *ADIC AML Administrator's Guide* for procedures to enter cartridges into the I/E/F bins.

Users must configure the Insert/Eject ports for the AML PVR using the configuration files `/var/hpss/etc/AML_EjectPort.conf` and `/var/hpss/etc/AML_InsertPort.conf`. The AML robot can have multiple insert and eject ports, which have the capability to handle different media types. These two configuration files in conjunction with the AMU AMS configuration files specify which Insert/Eject areas or bins the tapes should be placed in for insertion into the archive and the area to which they are ejected when the HPSS export command is used.

The AML PVR is supported by special bid only.

- **SCSI** - When importing new cartridges into HPSS, the cartridges must be entered into an input/output (I/O) slot before any HPSS import operations are performed. The HPSS SCSI PVR moves the cartridges into an available slot while importing into HPSS.
- **STK** - When importing new cartridges into HPSS, the cartridges must be entered into the STK robot before any HPSS import operations are performed. See the *STK Automated Cartridge System Library Software (ACSL) System Administrator's Guide* for procedures to enter cartridges.

Cartridges may be entered into the STK silo through various means, including:

- Using the enter command of ACSLS (StorageTek's robot software) and putting the cartridges into the Cartridge Access Port (CAP).
- Loading a cartridge into a CAP operating in ACSLS's "hot enter" mode.
- Bulk loading a robot and then performing an ACSLS audit of the slots that have been filled.
- STK supports cartridges without bar code external labels. These cartridges can be entered into the robot using the *venter* command at the ACSLS console. These cartridges are fully supported by HPSS, but the robot will eject the cartridges if they are scanned during an audit.

We strongly recommend that tape cartridges be imported into a robot that verifies external labels before mounting in a drive. This will ensure that a blank tape being imported has the appropriate external label before it is labeled. Once labeled, a cartridge can be moved to another PVR, such as an operator (hand) mounted PVR.

If labels are written on tape volumes prior to importing those volumes into HPSS, then care should be taken that the label has a format that HPSS can process. In particular, the first four characters of the OwnerID field should be "HPSS" and the last two characters of this fourteen character field should be "00". HPSS uses the last two characters to indicate the side number of the cartridge in anticipation of supporting cartridges that have multiple physical sides; for example, optical disk cartridges that contain an A and a B side. If these two characters are not set correctly, then the import will fail because the side information contained in the internal label will not match the side information passed to PVL in the Import request. If the start of the OwnerID field is not "HPSS", then the volume will be imported as a Foreign Label volume and the side will be set to zero.

When importing a non-removable disk volume into HPSS, the raw disk must already be defined in the host system. The import labels the volumes with HPSS labels.

The volumes can be imported into HPSS using the *Import Disk Volumes* window (*Section 8.1.1.3, "Import Disk Volumes window"* or the *Import Tape Volumes* window in the following section). Once

the volumes are imported, the HPSS storage space can be created using the *Create Disk Resources* or *Create Tape Resources* window.

The import step must be done for both disk and tape volumes. Ensure that the PVL, PVR, and appropriate Movers are up and running before initiating a requesting for the import operation. Also ensure that SSM is connected to the PVL.

When importing tapes, the maximum number of tape drives used by the import process is controlled by the **Max Number of Drives** field in the *Import Tape Volumes* window. Each tape must be mounted and its label must be written or checked. The PVL will attempt to evenly distribute the work across the specified number of tape drives. The administrator should consider the current and expected demand for tape drives of the type needed for the import, when setting this field.

When imports to an operator PVR are performed, requests to mount the tape cartridges are serialized. This is necessary to ensure proper labeling of unlabeled tapes. Any unlabeled cartridge mounted in a drive under control of the operator will be assumed to be the requested cartridge. Because any tape that is mounted will be labeled without machine checking of the external label, this operation is not without risk. It is strongly recommended that labeling of cartridges be performed by a robot that can verify the external label of a cartridge, if possible.

When SSM completes the import request, it reports the number of volumes imported in the window status field. Any error causes SSM to stop processing the request. If the reported number is less than the number requested, you should retry the import request after first addressing the reason for the initial failure. For the retry, the original list of volumes to import can be reused because the PVL will skip volumes that have already been imported.

8.1.1.1. Import Tape Volumes window

Import Tape Volumes

File Edit Help

PVR SCSI PVR LTO

Max Number of Drives

Manufacturer

Lot Number

☐ Defer Labeling

Media Type IBM - 3580 (LTO) Gen8 Tape

Import Type Default Import

Use list in external file to update Volume List

File Containing Volume List

Use starting label, count, and increment to update Volume List

Fill Count

Fill Increment

Volume Label

Volume List

Maximum Volumes Allowed

Total Count

Volume

Setting field: Media Type

This window allows the user to import tape volumes into the HPSS system, making them known to the PVL server. To make them known to the Core Server so they can be used by HPSS, storage resources must then be created for the volumes via the *Create Tape Resources* window.

When the **Max Number of Drives** is set to "1", the volumes are processed one at a time in sequence. If an error occurs, processing stops. The window completion message will report the number of successful imports, from which the volume causing the failure can be found. For example, if

you requested import of 100 volumes, and the completion message showed that 37 imports were successful, then the 38th volume in the **Volume List** caused the failure.

When the **Max Number of Drives** is greater than "1", the SSM System Manager effectively starts multiple import jobs (equal to **Max Number of Drives**) which run at the same time. Each job independently processes part of the **Volume List**, and stops when its part of the list is done or when an error occurs. The window completion message shows only the total of all successful imports from all of these jobs. It is therefore impossible to tell from this number which volume caused the failure. In this case it is important to review the *Alarms and Events* window to get more information on the failures.

If you request the import of a volume which was previously imported, the System Manager counts this as a successful import and goes on to the next volume in the list. This makes it easy to restart a partially completed import (after fixing the cause of the error which terminated the first request) by clicking the **Import** button again. There is no need to remove from the list the volumes which were imported successfully.

The list of the volumes to be imported may be constructed in any of three ways. Each volume name may be typed in one at a time, or a list of volume names may be automatically generated from a single entry, or a list of volume names may be specified from an input file. Volume names are not typed directly into the **Volume List** at the bottom of the window.

Any of the three entry methods may be repeated multiple times on the same window to add additional volumes to the list. All three entry methods or any combination of them may be used in succession on the same window.

To add a single volume name to the **Volume List**, set the **Fill Count** and the **Fill Increment** each to "1" and type the volume name into the **Volume Label** field. The volume name will be added to the **Volume List**.

To automatically generate a list of volume names in the **Volume List**, set the **Fill Count** to the desired number of volumes. Set the **Fill Increment** to the number by which each automatically generated label should differ from the previous one. Then type the starting volume name into the **Volume Label** field. The specified number of volume names will be added to the **Volume List**, each one larger than the previous entry by the specified **Fill Increment**.

Example:

```
Fill Count = 6
Fill Increment = 10
Volume Label = "AA0070"
```

Labels automatically inserted into Volume List:

```
"AA0070 "
"AA0080 "
"AA0090 "
"AA0100 "
"AA0110 "
"AA0120 "
```

When an addition produces overflow in a column, numerical columns are carried over properly to alphabetic columns and vice versa.

Example:

```
Fill Count = 6
Fill Increment = 2000
Volume Label= "AA7329"
```

Labels automatically inserted into Volume List:

```
"AA7329 "
"AA9329 "
"AB1329 "
"AB3329 "
"AB5329 "
"AB7329 "
```

The filling will not occur and an error will be displayed if the specified values would generate an invalid volume label (for example, one greater than zzz999).

To specify a list of volumes from a file, create a file containing the name of each volume to be imported on a separate line. Volume names must be six alphanumeric characters. No other characters are allowed on the line. The file must be accessible from the host on which the hpssgui is executing. Enter the name of the file in the **File Containing Volume List** field. The volume names from the file will be added to the **Volume List**.

Field descriptions.

PVR

The descriptive name of the PVR into which the tape will be imported.

Max Number of Drives

The maximum number of tape drives to use when running the tape imports. Valid values are any positive integer value, up to the number of drives available, to a maximum of 256. The default value is "1".

Manufacturer

The maker of the media to be imported. Any character string is valid. The field may be left blank.

Lot Number

The manufacturer's lot number for the media. Any character string is valid. The field may be left blank.

Defer Labeling

If selected, tape(s) in this import will not be labeled and will instead be labeled when they are first used for a write.

Media Type

The type of media to be imported.

Import Type

The type of import to be performed. Valid values are `Default Import`, `Scratch Import`, and `Overwrite Import`. The default value is `Default Import`. Specifying a `Scratch Import` type will usually cause an HPSS label to be written onto the media, potentially causing any data already on the media to be lost. Specifying an `Overwrite Import` type will cause the HPSS label to be rewritten, if an existing HPSS or Foreign label already exists, provided it is for the same Volume

ID. This allows tapes to be rewritten in newer drives (for example, tapes originally created on a 3590 drive to be reused on a 3590E drive). Based on the **Import Type**, the import request will be processed depending on how the media is currently labeled. See *Section 8.1.1.2, “Selecting import type for tape cartridges”* for more information on selecting the appropriate import type.

File Containing Volume List

The name of an external file (accessible from the host on which the **hpssgui** is executing) containing a list of volume labels to be added to the end of the **Volume List**.

Fill Count

The number of labels to be added to the **Volume List** when a value is typed into the **Volume Label** field. Valid values are positive integers greater than or equal to "1". The default value is "1". If the value specified for **Fill Count** would cause the list to grow beyond the Maximum Volumes Allowed, an error message is displayed (see **Maximum Volumes Allowed** below).

Fill Increment

This field determines how each new volume label is generated. The **Fill Increment** is the number by which each automatically generated label will differ from the previous one when a value is typed into the **Volume Label** field and the **Fill Count** is greater than "1". Valid values are any positive integer. The default value is "1". The **Fill Increment** is added to the least significant part of a volume label to generate a new label.

Volume Label

The six-character alpha-numeric label of a volume to be imported. The label is added to the end of the **Volume List**. If **Fill Count** is greater than "1", multiple labels are generated using the entered label as a starting point.

Maximum Volumes Allowed

The maximum number of volume labels that will fit in the **Volume List**. The value is 10,000 and is set by SSM. This field is non-editable.

Total Count

The total number of tapes to be imported. This is an informational field reflecting the number of volume names generated in the **Volume List** and is not directly editable.

Volume List

A list of volume labels specifying the volumes to be imported. You cannot enter labels directly into this list, but must construct it in one of the three ways described above. Use the scrollbar to move up and down the list of labels.

You can select one or more labels from the list for the purpose of deleting them from the list (see **Clear Selected** below). To select one label, click on it with the mouse; the selection will be highlighted. To select a range, select one label; then select a second label while holding down the Shift key. The selected labels, and all labels between the two, will be highlighted. You may also hold down the Control key while selecting, to select or deselect individual labels without affecting any of the others.

Buttons.**Clear All**

Clears the **Volume List**, and resets **Fill Count** and **Fill Increment** to "1".

Clear Selected

If one or more volumes are selected (highlighted) in the **Volume List**, clicking on this button will remove them from the list. Note that this does not actually delete anything from HPSS.

Import

Begins the volume import using the displayed data. A start message is displayed on the message line at the bottom of the window, and all window features are disabled, except for the **Dismiss** button; this prevents you from modifying any data on the window while the import is in progress.

It is a good idea to have the *Alarms and Events* window open while import proceeds, so that all relevant log messages can be viewed. For tape imports, it is also helpful to have the *Tape Mount Requests* window open, where you can see mount requests for each volume come and go as they are imported.

When the import is finished, a completion message is displayed and the window features are sensitized again. You may dismiss the window before completion; however, completion messages will be displayed in a pop-up window. At this point you can begin entering data for another import, or you can dismiss the window.

8.1.1.2. Selecting import type for tape cartridges

The following table lists information for selecting tape import types.



Any existing data on a volume imported into HPSS will be overwritten and subsequently lost.

Table 8.1. Tape import types

Current tape label	Default import	Overwrite import	Scratch import
A UniTree label (an ANSI label without a trailing tape mark) with a correct Volume ID	Tape imported	Label written, tape imported	Tape imported
An ANSI (non-UniTree) or HPSS label with a correct Volume ID (the Volume ID on the label is as expected by HPSS)	Tape imported	Label written, tape imported	Tape imported
An ANSI or HPSS label with an incorrect Volume ID (the Volume ID on the label is different from the Volume ID expected by HPSS)	Tape not imported	Tape not imported	Tape not imported
Random data (for example, a tar file)	Tape not imported	Label written, tape imported	Label written, tape imported
No data (two tapemarks at the start of tape)	Label written, tape imported	Label written, tape imported	Label written, tape imported

Current tape label	Default import	Overwrite import	Scratch import
Unreadable (for example, some brand new tapes or a degaussed tape; also possibly a tape written at a different density)	Tape not imported	Label written, tape imported	Label written, tape imported

HPSS always attempts to read a label at the beginning of a tape when performing an import. Some tape drives and device drivers will report errors when asked to read unreadable tapes (as described above). If this happens, manually write two tape marks at the start of the tape and retry the import. Most UNIX systems provide the `mt` command, which can be used to write tape marks.

An HPSS label is basically just an ANSI label. A few characters are changed to identify the tape as having been labeled by HPSS. Both label types are supported by HPSS; tapes already beginning with an ANSI label are not relabeled unless the overwrite import option is selected and volume labels match.

8.1.1.3. Import Disk Volumes window

This window allows the user to import disk volumes into the HPSS system, making them known to the PVL and PVR servers. To make them known to the Core Server so they can be used, they must be created via the Core Server's *Create Disk Resources* window.

The SSM System Manager processes the volumes one at a time in sequence. If it encounters an error, it stops and returns. The window completion message will report the number of successful imports, from which the volume causing the failure can be found. For example, if you requested import of 100 volumes, and the completion message showed that 37 imports were successful, then the 38th volume in the **Volume List** caused the failure.

If you request the import of a volume which was previously imported, the System Manager counts this as a successful import and goes on to the next volume. This makes it easy to restart a partially completed import (after fixing the cause of the error which terminated the first request) by clicking the **Import** button again. There is no need to remove the already imported volumes from the list.

The list of the volumes to be imported may be constructed in any of three ways. Each volume name may be typed in one at a time, or a list of volume names may be automatically generated from a single entry, or a list of volume names may be specified from an input file. Volume names are not typed directly into the **Volume List** at the bottom of the window.

Any of the three entry methods may be repeated multiple times on the same window to add additional volumes to the list. All three entry methods or any combination of them may be used in succession on the same window.

To add a single volume name to the **Volume List**, set the **Fill Count** and the **Fill Increment** each to "1" and type the volume name into the **Volume Label** field. The volume name will be added to the **Volume List**.

To automatically generate a list of volume names in the **Volume List**, set the **Fill Count** to the desired number of volumes. Set the **Fill Increment** to the number by which each automatically generated label should differ from the previous one. Then type the starting volume name into the **Volume Label** field. The specified number of volume names will be added to the **Volume List**, each one larger than the previous entry by the specified **Fill Increment**.

Example:

```
Fill Count = 6
Fill Increment = 10
Volume Label = "AA0070"
```

Labels automatically inserted into Volume List:

```
"AA0070"
"AA0080"
"AA0090"
"AA0100"
"AA0110"
"AA0120"
```

When an addition produces overflow in a column, numerical columns are carried over properly to alphabetic columns and vice versa.

Example:

```
Fill Count = 6
Fill Increment = 2000
Volume Label= "AA7329"
```

Labels automatically inserted into Volume List:

```
"AA7329 "  
"AA9329 "  
"AB1329 "  
"AB3329 "  
"AB5329 "  
"AB7329 "
```

The filling will not occur and an error will be displayed if the specified values would generate an invalid volume label (for example, one greater than zzz999).

To specify a list of volumes from a file, create a file containing the name of each volume to be imported on a separate line. Volume names must be six alphanumeric characters. No other characters are allowed on the line. The file must be accessible from the host on which the **hpssgui** is executing. Enter the name of the file in the **File Containing Volume List** field. The volume names from the file will be added to the **Volume List**.

Field descriptions.

Media Type

The type of media to be imported.

Import Type

The type of import to be performed. Valid values are `Default Import`, `Scratch Import`, and `Overwrite Import`. Based on the import type, the import request will be processed depending on how the media is currently labeled. See *Section 8.1.1.4, "Selecting import type for disk volumes"* for more information on selecting the appropriate import type.

File Containing Volume List

The name of an external file containing a list of volume labels to be added to the end of the **Volume List**.

Fill Count

The number of labels to be added to the **Volume List** when a value is typed into the **Volume Label** field. Valid values are positive integers greater than or equal to "1". The default value is "1". If the value specified for **Fill Count** would cause the list to grow beyond the **Maximum Volumes Allowed**, an error message is displayed (see **Maximum Volumes Allowed** below).

Fill Increment

This field determines how each new volume label is generated. The **Fill Increment** is the number by which each automatically generated label will differ from the previous one when a value is typed into the **Volume Label** field and the **Fill Count** is greater than "1". Valid values are any positive integer. The default value is "1". The **Fill Increment** is added to the least significant part of a volume label to generate a new label.

Volume Label

The six-character alpha-numeric label of a volume to be imported. The label is added to the end of the **Volume List**. If **Fill Count** is greater than "1", multiple labels are generated using the entered label as a starting point.

Maximum Volumes Allowed

The maximum number of volume labels that will fit in the **Volume List**. The value is 10,000 and is set by SSM. This field is non-editable.

Total Count

The total number of disks to be imported. This is an informational field reflecting the number of volume names generated in the **Volume List** and is not directly editable.

Volume List

A list of volume labels specifying the volumes to be imported. You cannot enter labels directly into this list, but must construct it using one of the three ways described above. Use the scrollbar to move up and down the list of labels.

You can select one or more labels from the list for the purpose of deleting them from the list (see **Clear Selected** below). To select one label, click on it with the mouse; the selection will be highlighted. To select a range, select one label; then select a second label while holding down the Shift key. The selected labels, and all labels between the two, will be highlighted. You may also hold down the Control key while selecting, to select or deselect individual labels without affecting any of the others.

Buttons.**Clear All**

Clears the **Volume List**, and resets **Fill Count** and **Fill Increment** to "1".

Clear Selected

If one or more volumes are selected (highlighted) in the **Volume List**, clicking on this button will remove them from the list. Note that this does not actually delete anything from HPSS.

Import

Begins the volume import using the displayed data. A start message is displayed on the message line at the bottom of the window, and all window features are disabled, except for the **Dismiss** button; this prevents you from modifying any data on the window while the import is in progress.

It is a good idea to have the *Alarms and Events* window open while import proceeds, so that all relevant log messages can be viewed.

When the import is finished, a completion message is displayed and the window features are sensitized again. You may dismiss the window before completion; however, completion messages will be displayed in a pop-up window. At this point you can begin entering data for another import, or you can dismiss the window.

8.1.1.4. Selecting import type for disk volumes

The following table lists information for selecting disk import types.

Table 8.2. Disk import types

Current disk label	Default import	Overwrite import	Scratch import
An ANSI or HPSS label with a correct Volume ID (the Volume ID on the label is as expected by HPSS)	Disk imported	Label written, disk imported	Disk imported

Current disk label	Default import	Overwrite import	Scratch import
An ANSI or HPSS label with an incorrect Volume ID (the Volume ID on the label is different from the Volume ID expected by HPSS)	Disk not imported	Disk not imported	Disk not imported
No label	Label written, disk imported	Label written, disk imported	Label written, disk imported

8.1.2. Creating storage resources

Creating storage resources is the process of creating Core Server metadata to manage the space on HPSS volumes. The volumes must be imported into HPSS (see *Section 8.1.1, “Importing volumes into HPSS”*) before resources may be created on them.

Before storage resources can be created, the administrator must determine which storage subsystem will receive the new storage, and which storage class they will be placed in. The choice of storage subsystem determines the Core Server that will be used.

Ensure that the PVL and the appropriate Core Server are running before initiating a request to create storage resources. Also ensure that SSM is connected to the appropriate Core Server.

The selected Core Server will create a number of HPSS virtual volumes. All new tape virtual volumes will have identical characteristics (including type, estimated size, stripe size, stripe width, block size, and blocks between tape marks). When creating disk storage resources, the administrator may override the default PV Size and enter a value specific to the circumstance.

Virtual volumes are created by the Core Server by grouping physical volumes together in groups of N, where N is the total stripe width of the resulting virtual volume. The physical volumes are arranged into virtual volumes according to the table in the resource creation window, which the administrator should modify as desired before clicking the **Create** button. Every physical volume is placed into a virtual volume before it is usable by HPSS, even if there is only one physical volume per virtual volume.

It is important to understand that the storage class characteristics in effect when volumes are created are assigned to the new volumes. Once volumes are created, changes to the storage class have no effect on the characteristics of existing volumes. The storage class characteristics (such as block sizes and stripe widths) are copied to the volumes' metadata and become a permanent part of the definition of the volumes.

The Core Server creates the necessary metadata structures for each of the new virtual volumes. Each new virtual volume is immediately available for use.

Note that new tape resources are not assigned to a file family. Tapes are assigned to file families from the unused tapes in the storage class as they are needed to satisfy requests to write to tape in a given family. Tapes cannot be pre-assigned to file families.

8.1.2.1. Create Tape Resources window

Create Tape Resources

File Edit Help

Initialization Fields – Fill These In First

Storage Class **2 | Storage Class 2**

Core Server **Core Server**

Optional or Informational Fields

PVs in Each VV **1** (DSW + PSW)

Fill Direction **Horizontal**

PV Size **200MB**

RAIT Options

Data Stripe Width (DSW) **1**

Parity Stripe Width (PSW) **0** (0 = RAIT is disabled)

Minimum Write Parity **0**

Number of RAIT Engines **0**

☐ Read Verification ☐ Mount Minimum PVs

Use list in external file to update Volume List

File Containing Volume List

Use starting label, count, and increment to update Volume List

Fill Count **1**

Fill Increment **1**

Volume Label

Volume List

Total Count **0**

VVs To Create **0**

Volume

0/0

Clear All Create Resources Dismiss

This window is used to create tape storage resources, tape virtual volumes, in a Core Server. The tapes must first be imported to the appropriate PVL.

The names of the volumes may be entered into the window one at a time, or a list of volume names may be automatically generated from a single entry. Volume names are not entered directly into the **Volume List** at the bottom of the window but are typed into the **Volume Label** field.

To add a single volume name to the **Volume List**, set the **Fill Count** and the **Fill Increment** each to "1" and type the volume name into the **Volume Label** field. The volume name will be added to the **Volume List**.

To automatically generate a list of volume names in the **Volume List**, set the **Fill Count** to the desired number of volumes. Set the **Fill Increment** to the number by which each automatically generated label should differ from the previous one. Then type the starting volume name into the **Volume Label** field. The specified number of volume names will be added to the **Volume List**, each one larger than the previous entry by the specified **Fill Increment**.

Example:

```
Fill Count = 6
Fill Increment = 10
Volume Label = "AA0070"
```

Labels automatically inserted into Volume List table:

```
"AA0070 "
"AA0080 "
"AA0090 "
"AA0100 "
"AA0110 "
"AA0120 "
```

When an addition produces overflow in a column, numerical columns are carried over properly to alphabetic columns and vice versa.

Example:

```
Fill Count = 6
Fill Increment = 2000
Volume Label= "AA7329"
```

Labels automatically inserted into Volume List table:

```
"AA7329 "
"AA9329 "
"AB1329 "
"AB3329 "
"AB5329 "
"AB7329 "
```

Once the Core Server completes the create request, SSM reports the number of physical volumes and virtual volumes created in the window status field.

Field descriptions.

Initialization Fields - Fill These In First

Storage Class

The name of the storage class whose characteristics will be used to create the requested virtual volumes. The stripe width of this storage class will determine the number of columns in the **Volume List** table at the bottom of the window.

Core Server

The name of the Core Server that will process the create request and in which the tape storage resources will be created. The **Storage Class** field must be filled in before selecting a Core Server.

Optional or Informational Fields**PVs in Each VV**

The **Data Stripe Width** plus **Parity Stripe Width** of the selected storage class. This field may not be changed.

Fill Direction

Determines which direction filling takes place in the table. By default, **Fill Direction** is **Horizontal**. Select the desired value for this field before entering a **Volume Label**. **Fill Direction** will not be enabled if the **PVs in Each VV** is "1".

PV Size

The size, in bytes, of each physical volume used. The default value is the value from the storage class. For tapes, this is an estimate and is informational only.

RAIT options**Data Stripe Width**

The number of PVs in each new VV that will contain file information. This field may not be changed.

Parity Stripe Width

The number of PVs in each new VV that will contain RAIT parity information. When Parity Stripe Width is greater than zero, the VV is RAIT. When zero, the VV is not RAIT. This field may not be changed.

Minimum Write Parity

The minimum number of parity PVs that must remain writable in order for the RAIT VV to be writable. This field may not be changed.

Number of RAIT Engines

The number of RAIT Engines that will be used when reading or writing the VVs. This field may not be changed.

Read Verification

When set, parity information in the RAIT VV will be used to verify the integrity of data read from the VV. This field can be set or cleared before the VVs are created.

Mount Minimum PVs

When set, the smallest number of PVs necessary will be mounted when reading the RAIT VV. When **Read Verification** is not set, this will be equal to the value of DSW; and when **Read Verification** is set, this will equal (DSW + 1).

Use list in external file to update Volume List**File Containing Volume List**

The name of an external file (accessible from the host on which the **hpssgui** is executing) containing a list of volume labels to be added to the end of the **Volume List**.

Use starting label, count, and increment to update Volume List

Fill Count

The number of labels to be added to the **Volume List** when a value is typed into the **Volume Label** field. Valid values are positive integers greater than or equal to "1" up to the number of drives available. The default value is "1". If the list fills before the **Fill Count** is exhausted, filling stops and a message is displayed.

Fill Increment

This field determines how each new volume label is generated. The **Fill Increment** is the number by which each automatically generated label will differ from the previous one when a value is in the **Volume Label** field and the **Fill Count** is greater than "1". Valid values are positive integers up to the number of available drives. The default value is "1". A volume label is six alphanumeric characters. The **Fill Increment** is added to the least significant portion of a volume label to generate a new label.

Volume Label

The next label that will be generated and inserted into the table. The labels are generated when the field loses focus. Select **Fill Count**, **Fill Increment**, and **Fill Direction** before entering a value in this field.

Volume List

A list of volume labels specifying the volumes used to create resources. You cannot enter labels directly into this list but must instead use the **Fill Count**, **Fill Increment**, and **Volume Label** fields. Use the scrollbar to navigate the list of labels.

Total Count

The total number of physical volumes in the generated Volume List. This is an informational field and not directly enterable.

VVs To Create

The number of virtual volumes that will be created. This is an informational field and not directly enterable.

Buttons.

Clear All

This button clears the volume list and **Volume Label** field.

Create Resources

Begins the process of creating HPSS storage resources using the displayed data. A start message is displayed on the status bar, and all window features are disabled, except for the **Dismiss** button. This prevents the user from modifying any data on the window while the create request is in progress.

8.1.2.2. Create Disk Resources window

Create Disk Resources

File Edit Help

Initialization Fields - Fill These In First

Storage Class 1000 | Storage Class 1000

Core Server Core Server

Optional or Informational Fields

PVs in Each VV 1

Fill Direction Horizontal

PV Size 42GB

☒ Create DOWN

Use list in external file to update Volume List

File Containing Volume List

Use starting label, count, and increment to update Volume List

Fill Count 1

Fill Increment 1

Volume Label DISK42

Volume List

Total Count 1

VVs To Create 1

Volume
DISK4200

1/1

Clear All Create Resources Dismiss

Setting field: Volume Label

This window is used to create disk storage resources, disk virtual volumes, in a Core Server. The disks must first be imported to the appropriate PVL.

The names of the volumes may be entered into the window one at a time, or a list of volume names may be automatically generated from a single entry. Volume names are not entered directly into the **Volume List** at the bottom of the window but are typed into the **Volume Label** field.

To add a single volume name to the **Volume List**, set the **Fill Count** and the **Fill Increment** each to "1" and type the volume name into the **Volume Label** field. The volume name will be added to the **Volume List**.

To automatically generate a list of volume names in the **Volume List**, set the **Fill Count** to the desired number of volumes. Set the **Fill Increment** to the number by which each automatically generated label should differ from the previous one. Then type the starting volume name into the **Volume Label** field. The specified number of volume names will be added to the **Volume List**, each one larger than the previous entry by the specified **Fill Increment**.

Example:

```
Fill Count = 6
Fill Increment = 10
Volume Label = "AA0070"
```

Labels automatically inserted into Volume List table:

```
"AA0070 "
"AA0080 "
"AA0090 "
"AA0100 "
"AA0110 "
"AA0120 "
```

When an addition produces overflow in a column, numerical columns are carried over properly to alphabetic columns and vice versa.

Example:

```
Fill Count = 6
Fill Increment = 2000
Volume Label= "AA7329"
```

Labels automatically inserted into Volume List table:

```
"AA7329 "
"AA9329 "
"AB1329 "
"AB3329 "
"AB5329 "
"AB7329 "
```

Once the Core Server completes the create request, SSM reports the number of physical volumes and virtual volumes created in the window status field.

Field descriptions.

Initialization Fields - Fill These In First

Storage Class

The name of the storage class whose characteristics will be used to create the requested virtual volumes. The stripe width of this storage class will determine the number of columns in the **Volume List** table at the bottom of the window.

Core Server

The name of the Core Server that will process the create request and in which the disk storage resources will be created. The **Storage Class** field must be filled in before selecting a Core Server.

Optional or Informational Fields**PVs in Each VV**

The stripe width of the selected storage class. This field may not be changed.

Fill Direction

Determines which direction filling takes place in the table. By default, **Fill Direction** is **Horizontal**. Select the desired value for this field before entering a **Volume Label**. **Fill Direction** will not be enabled if the **PVs in Each VV** is "1".

PV Size

The size, in bytes, of each physical volume used. The default value is the value from the storage class. For disks, it should be changed on this window to the value that matches the actual value of each disk being imported.

Create DOWN

Creates the disk volumes in the DOWN state. DOWN disks will not be immediately usable and may be brought online by changing the volumes' VV Conditions to another state (RO, RW, or RWC). See *Section 8.5.4.1, "Core Server Disk Volume information window"* for information about a disk volume's VV Condition.

Use list in external file to update Volume List**File Containing Volume List**

The name of an external file (accessible from the host on which the **hpssgui** is executing) containing a list of volume labels to be added to the end of the **Volume List**.

Use starting label, count, and increment to update Volume List**Fill Count**

The number of labels to be added to the **Volume List** when a value is typed into the **Volume Label** field. Valid values are positive integers greater than or equal to "1" up to the number of drives available. The default value is "1". If the list fills before the **Fill Count** is exhausted, filling stops and a message is displayed.

Fill Increment

This field determines how each new volume label is generated. The **Fill Increment** is the number by which each automatically generated label will differ from the previous one when a value is in the **Volume Label** field and the **Fill Count** is greater than "1". Valid values are positive integers up to the number of available drives. The default value is "1". A volume label is six alphanumeric characters. The **Fill Increment** is added to the least significant portion of a volume label to generate a new label.

Volume Label

The next label that will be generated and inserted into the table. The labels are generated when the field loses focus. Select **Fill Count**, **Fill Increment**, and **Fill Direction** before entering a value in this field.

Volume List

A list of volume labels specifying the volumes used to create resources. You cannot enter labels directly into this list but must instead use the **Fill Count**, **Fill Increment**, and **Volume Label** fields. Use the scrollbar to navigate the list of labels.

Total Count

The total number of physical volumes in the generated Volume List. This is an informational field and not directly enterable.

VVs To Create

The number of virtual volumes that will be created. This is an informational field and not directly enterable.

Buttons.**Clear All**

This button clears the volume list and **Volume Label** field.

Create Resources

Begins the process of creating HPSS storage resources using the displayed data. A start message is displayed on the status bar, and all window features are disabled, except for the **Dismiss** button. This prevents the user from modifying any data on the window while the create request is in progress.

8.2. Removing storage space

Removing storage space from the HPSS system is a two-step process:

1. **Delete Resources.** In the first step, the Core Server virtual volume metadata is deleted; this process is called deleting storage resources.
2. **Export Volumes.** In the second step, the physical volumes that belonged to the virtual volumes are exported from the PVL.

The second step, exporting the volumes from the PVL, is not always necessary. Once the resources are deleted in Step 1, new storage resources can be created on the volumes without the need to export the volumes from the system and reimport them. Exporting the volumes is necessary only if the volumes are not going to be reused in the HPSS system, or if for some reason they must be relabeled. The site may wish to export the volumes if they are to be imported into a different PVR in the system, but this function can also be accomplished with the Move Cartridges operation.

8.2.1. Deleting storage resources

Storage resources can be removed from HPSS volumes by the use of the *Delete Resources* window or the **remove** utility.

Once resources have been deleted, the volumes may be reused or removed from the HPSS system. To reuse the volumes, create new storage resources on them using the *Create Resources* window. To

remove them entirely from the HPSS system, export them from the PVL using the *Export Volumes* window.

8.2.1.1. Rules for deleting resources

Volumes on which resources are to be deleted must be empty. There must be no storage segments associated with the volumes. The *Core Server Disk Volume* or *Core Server Tape Volume* window must show a zero in the **Number of Active Segments** field, and the VV Condition must be EMPTY. If any storage segments remain on a volume, the resource deletion operation will fail. You can use the volume **repack** features of HPSS to empty disk and tape volumes so that their resources may be deleted.

The Delete Resources operation may be performed on either retired or non-retired volumes.

Successful deletions result in the removal of the Core Server's metadata records from the appropriate subsystem. The volume is then no longer usable for storage by HPSS, but it is still recognized by the PVL.

8.2.1.2. Delete Resources window

The screenshot shows the 'Delete Resources' window. It has a title bar with the text 'Delete Resources' and standard window controls. Below the title bar are menu items: 'File', 'Edit', and 'Help'. The main area is divided into several sections. The first section is titled 'Use list in external file to update Volume List' and contains a text input field labeled 'File Containing Volume List'. The second section is titled 'Use starting label, count, and increment to update Volume List' and contains three input fields: 'Fill Count' with the value '1', 'Fill Increment' with the value '1', and 'Volume Label'. Below these is a section titled 'Volume List' which contains two input fields: 'Maximum Volumes Allowed' with the value '10000' and 'Total Count' with the value '0'. Below these fields is a list box with a single header 'Volume' and a scrollbar. The list box shows '0/0' items. At the bottom of the window are four buttons: 'Clear All', 'Clear Selected', 'Delete Resources', and 'Dismiss'.

This window is used to delete empty or unused storage resources (that is, virtual volumes) from a Core Server. The volumes must have no storage segments on them when they are removed.

The list of the volumes to be deleted may be constructed in any of three ways. Each volume name may be typed in one at a time, or a list of volume names may be automatically generated from a single

entry, or a list of volume names may be specified from an input file. Volume names are not typed directly into the **Volume List** at the bottom of the window.

Any of the three entry methods may be repeated multiple times on the same window to add additional volumes to the list. All three entry methods or any combination of them may be used in succession on the same window.

To add a single volume name to the **Volume List**, set the **Fill Count** and the **Fill Increment** each to "1" and type the volume name into the **Volume Label** field. The volume name will be added to the **Volume List**.

To automatically generate a list of volume names in the **Volume List**, set the **Fill Count** to the desired number of volumes. Set the **Fill Increment** to the number by which each automatically generated label should differ from the previous one. Then type the starting volume name into the **Volume Label** field. The specified number of volume names will be added to the **Volume List**, each one larger than the previous entry by the specified **Fill Increment**.

Example:

```
Fill Count = 6
Fill Increment = 10
Volume Label = "AA0070"

Labels automatically inserted into Volume List:
"AA0070"
"AA0080"
"AA0090"
"AA0100"
"AA0110"
"AA0120"
```

When an addition produces overflow in a column, numerical columns are carried over properly to alphabetic columns and vice versa.

Example:

```
Fill Count = 6
Fill Increment = 2000
Volume Label= "AA7329"

Labels automatically inserted into Volume List:
"AA7329"
"AA9329"
"AB1329"
"AB3329"
"AB5329"
"AB7329"
```

The filling will not occur and an error will be displayed if the specified values would generate an invalid volume label (for example, one greater than zzz999).

To specify a list of volumes from a file, create a file containing the name of each volume to be deleted on a separate line. Volume names must be six alphanumeric characters. No other characters are allowed on the line. The file must be accessible from the host on which the hpssgui is executing. Enter

the name of the file in the **File Containing Volume List** field. The volume names from the file will be added to the **Volume List**.

Field descriptions.

File Containing Volume List

The name of an external file containing a list of volume labels to be added to the end of the **Volume List**.

Fill Count

The number of volume labels to be added to the end of the **Volume List** when the **Volume Label** field is next modified.

Fill Increment

In a multiple-label fill, this field determines how each new PV label is generated from the previous one. Each new volume label entered in the table is the previous entry incremented by this amount.

Volume Label

The label to be added to the end of the PV list. If **Fill Count** is greater than "1", this will be the first label of a fill sequence. The label must be exactly six characters long. Entering a value in this field causes the table entries to be generated.

Maximum Volumes Allowed

The maximum number of volume labels that will fit in the **Volume List**. This field is informational and not enterable.

Total Count

The total number of volumes in the list to be deleted. This is the current size of the generated **Volume List**.

Volume List

A list of volume labels of volumes to be deleted. You cannot enter labels directly into this list but must construct it using one of the three methods described above. You can select one or more labels from the list for the purpose of removing them from the list (see **Clear Selected** below). To select one label, click on it with the mouse; the selection will be highlighted. To select a range, select one label; then select a second label while holding the Shift key. The selected labels, and all labels between them, will be highlighted. You may also hold down the Control key while selecting, to select or deselect individual labels without affecting any other selections.

Buttons.

Clear All

Clears the **Volume List** and resets **Fill Count** and **Fill Increment**.

Clear Selected

If one or more volumes are selected (highlighted) in the **Volume List**, clicking on this button will remove them from the list. Note that this does not delete anything from HPSS.

Delete Resources

Begins the deletion of resources using the displayed data. A start message is displayed on the status bar, and all window features are disabled except for the **Dismiss** button. This prevents the user from modifying any data on the window while the delete request is in progress.

Related information.

Section 8.4, “Dealing with a space shortage”.

8.2.2. Exporting volumes from HPSS

Exporting volumes from HPSS is the process of removing tape cartridges or disk volume from HPSS control.

8.2.2.1. Rules for exporting volumes

Tape cartridges may be physically exported from any managing robotic library. To export a tape cartridge from HPSS, the administrator must be familiar with the operation of the PVR from which the cartridge will be removed because the physical cartridge ejection process differs among the PVRs supported by HPSS:

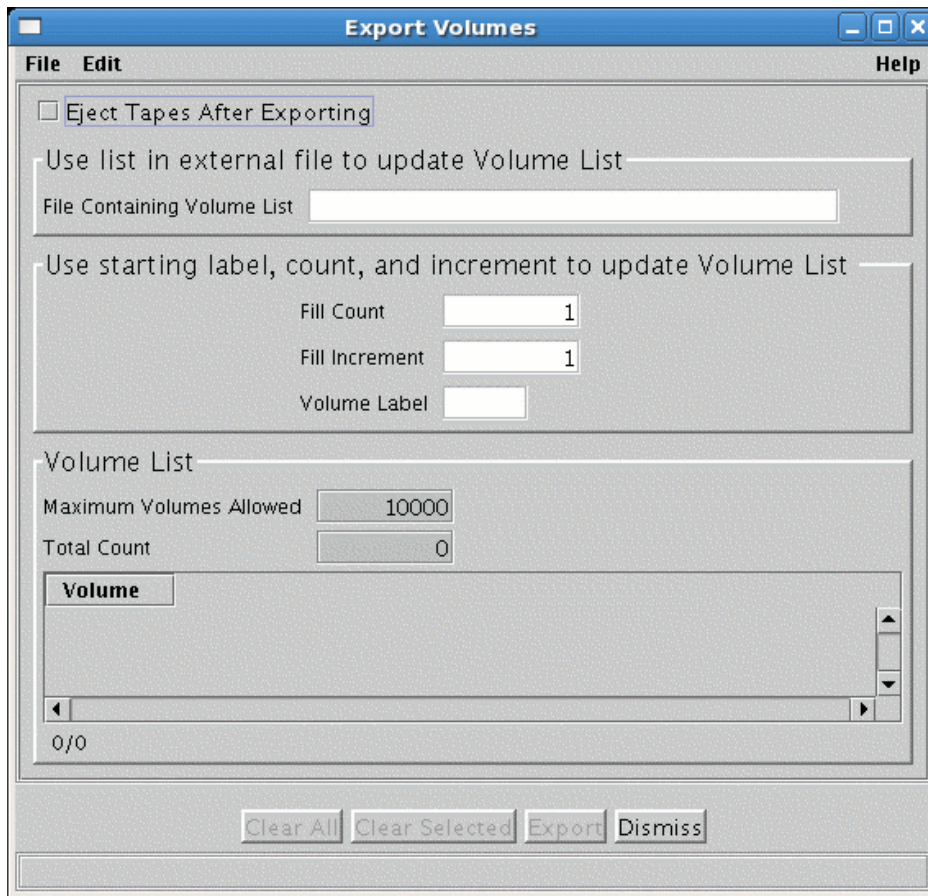
- **Operator** - No cartridge ejection step is necessary (or possible).
- **AML** - The exported cartridges will be placed in the E/I/F ports and can then be manually removed.
- **SCSI** - When an HPSS tape export command is issued, the cartridges will be placed in the input/output (I/O) slot and can then be manually removed.
- **STK** - The exported cartridges will be placed in the CAP and can then be manually removed. The export operation is *not* completed until cartridges are manually removed.

Only imported but unallocated volumes can be exported. Volumes allocated to a Core Server can be unallocated by deleting all Core Server metadata that describe the volume using the procedure described in Section 8.2.1, “Deleting storage resources”.

When exporting a removable cartridge from a robot, the cartridges will normally be sent to the convenience I/O port if it is available (this will not be done if the Export Without Eject option is used). If no locations are available for the cartridge to be ejected, an alarm will appear on the *Alarms and Events* window and HPSS will periodically retry the eject operation.

Use the *Export Volumes* window in the following section to perform the actual export operation.

8.2.2.2. Export Volumes window



This window allows you to export tape and disk volumes from the HPSS system. Exporting a volume is equivalent to telling HPSS that the volume no longer exists. Before volumes can be exported, the Core Server storage resources that describe the volumes must be deleted using the procedure described in *Section 8.2.1, “Deleting storage resources”*.

The list of the volumes to be exported may be constructed in any of three ways. Each volume name may be typed in one at a time, or a list of volume names may be automatically generated from a single entry, or a list of volume names may be specified from an input file. Volume names are not entered directly into the **Volume List** at the bottom of the window.

Any of the three entry methods may be repeated multiple times on the same window to add additional volumes to the list. All three entry methods or any combination of them may be used in succession on the same window.

To add a single volume name to the **Volume List**, set the **Fill Count** and the **Fill Increment** fields each to "1" and type the volume name into the **Volume Label** field. The volume name will be added to the **Volume List**.

To automatically generate a list of volume names in the **Volume List**, set the **Fill Count** to the desired number of volumes. Set the **Fill Increment** to the number by which each automatically generated label should differ from the previous one. Then type the starting volume name into the **Volume Label** field. The specified number of volume names will be added to the **Volume List**, each one larger than the previous entry by the specified **Fill Increment**.

Example:

```
Fill Count = 6
Fill Increment = 10
Volume Label = "AA0070"

Labels automatically inserted into Volume List:
"AA0070 "
"AA0080 "
"AA0090 "
"AA0100 "
"AA0110 "
"AA0120 "
```

When an addition produces overflow in a column, numerical columns are carried over properly to alphabetic columns and vice versa.

Example:

```
Fill Count = 6
Fill Increment = 2000
Volume Label= "AA7329"

Labels automatically inserted into Volume List:
"AA7329 "
"AA9329 "
"AB1329 "
"AB3329 "
"AB5329 "
"AB7329 "
```

The filling will not occur and an error will be displayed if the specified values would generate an invalid volume label (for example, one greater than zzz999).

To specify a list of volumes from a file, create a file containing the name of each volume to be exported on a separate line. Volume names must be six alphanumeric characters. No other characters are allowed on the line. The file must be accessible from the host on which the **hpssgui** is executing. Enter the name of the file in the **File Containing Volume List** field. The volume names from the file will be added to the **Volume List**.

The entire list of volumes is sent to the System Manager as a single export request. The System Manager processes them one at a time, and if it encounters an error, it stops and returns. In this case, the window status line will report the number of successful exports, and the volume name that caused the failure. Since the System Manager does not ignore volumes that have already been exported, but instead considers them failed exports, you cannot retry an export by resubmitting the same list of volumes. You will first have to remove the successfully exported volumes from the list. The **Clear Selected** button can be useful for this.

Once the export is completed, cartridges residing in robotic tape systems will be ejected by the controlling PVR (if **Eject Tapes after Exporting** is selected) and should be removed from the access port of the robot.

Field descriptions.**Eject Tapes After Exporting**

If this check box is selected, the exported tape volumes will also be ejected from the PVR.

File Containing Volume List

The name of an external file containing a list of volume labels to be added to the end of the **Volume List**.

Fill Count

The number of volume labels to be added to the end of the list when the **Volume Label** field is next modified. This number may be one or greater. If the list fills up before the **Fill Count** is exhausted, filling stops, and a message box is displayed (see **Maximum Volumes Allowed** below).

Fill Increment

In a multiple-label fill, this field determines how each new volume label is generated from the previous one. A volume label is six alphanumeric characters. The **Fill Increment** is added to the numerical part of a volume label to generate a new label.

Volume Label

The volume label to be added to the end of the **Volume List**. If **Fill Count** is greater than "1", this will be the first label of a fill sequence. The label must be exactly six characters long. Lower case characters entered into this field will be converted to upper case.

Maximum Volumes Allowed

The maximum number of volume labels that will fit in the **Volume List**. This value is 10,000 and is maintained by SSM. This field is non-editable.

Total Count

The total number of volumes currently in the **Volume List**.

Volume List

A list of volume labels specifying the volumes to be exported. You cannot enter labels directly into this list, but must construct it using one of the three ways described above. Use the scrollbar to move up and down the list of labels.

You can select one or more labels from the list for the purpose of deleting them from the list (see **Clear Selected** below).

Buttons.**Clear All**

Clears the **Volume List** and resets **Fill Count** and **Fill Increment** to "1".

Clear Selected

If one or more volumes are selected (highlighted) in the **Volume List**, clicking on this button will remove them from the list. Note that this does not actually delete anything from HPSS.

Export

Begins the volume export using the displayed data. A start message is displayed on the status bar at the bottom of the window, and all window features are disabled, except for the **Dismiss** button; this prevents you from modifying any data on the window while the export is in progress.

8.3. Monitoring storage space

One of the most important administrative tasks is to monitor storage space usage in the HPSS system. The goal is early detection of any potential space shortages so that corrective action can be taken in a timely manner. Each storage class is configured with a warning threshold and a critical threshold to inform SSM users when the free space in a storage class runs low. Warning and major alarms are sent to the *Alarms and Events* window periodically when the warning and critical thresholds are exceeded. Migration policies can be configured to automatically start migration when either threshold is exceeded.

When a storage class experiences a threshold-exceeded condition, an administrator may need to review the associated migration and purge policies for the storage class as well as the total storage space available to the storage class. The migration and purge policies may need to be modified to free up more space or to free up the space more frequently. In addition, the total storage space for the storage class may need to be reviewed to determine whether it is sufficient to accommodate the actual usage of the storage class.

Storage space in HPSS is monitored from the *Active Storage Classes* window.

8.3.1. Active Storage Classes window

The screenshot shows the 'Active Storage Classes' window. It features a menu bar (File, Edit, Column View) and a title bar. Below the title bar, there are two time-related fields: 'Time Updated by System Manager' and 'Time Received by Client', both showing 'Sep 25, 2008 7:19:49 AM'. The main area is a table with the following columns: ID, Name, Type, Thresh, % Used, Total Space, and Migr S. The table contains 16 rows of data. The right sidebar has a 'Help' button and a list of controls: Migration Controls, Purge Controls, Repack Volumes, Reclaim Volumes, Information, Storage Class Info, Configuration, Migration Policy, Purge Policy, Preferences, Edit, and Default. At the bottom, there are buttons for Freeze, Refresh, and Dismiss.

ID	Name	Type	Thresh	% Used	Total Space	Migr S
1	lwd - lv1 1 (classic)	Disk	OK	33	131,072 MB	Running
2	lwd - lv1 2 (classic)	Disk	OK	1	131,072 MB	Waiting
3	lwd (max)	Disk	OK	1	98,304 MB	Waiting
4	lwd - (var length)	Disk	OK	1	98,304 MB	Waiting
5	4wd (classic)	Disk	OK	1	262,144 MB	Waiting
6	lwd - Marc Test	Disk	OK	1	8,192 MB	None
9	lwd (no migr)	Disk	OK	0	0 MB	None
101	lwt (migr/purge)	Tape	OK	34	3 W	Waiting
102	2wt (no mig)	Tape	Warning	50	2 W	None
103	lwt (1st copy)	Tape	OK	34	3 W	Running
104	lwt (2nd copy)	Tape	OK	34	3 W	Waiting
105	lwt (no mig)	Tape	Warning	67	3 W	None
106	lwt (3580 Gen3)	Tape	Warning	50	2 W	None
107	lwt (9840D)	Tape	OK	0	0 W	None
108	lwt (T10K-A)	Tape	Critical	100	1 W	None
109	lwt (T10K-B)	Tape	OK	0	0 W	None

The *Active Storage Classes* window allows you to view a list of all active storage classes. An active storage class is defined as a storage class configured in HPSS to which storage volumes have been assigned. Volumes are assigned to a storage class by the Create Resources operation.

The information displayed in this window is obtained from the MPS, which obtains part of it from the Core Servers.



The Migration/Purge Server must be running and connected to SSM or the Active Storage Classes window will be stale. All Core Servers must be running or the storage classes managed by them will not be reported in this window.

If multiple storage subsystems exist, this window will display a list of storage classes for each subsystem. To display this window, from the *HPSS Health and Status* window, select the **Monitor** menu, and from that select the **Storage Classes Active** menu item.

The function buttons to the right of this list require that one storage class row be selected from the list. To select a row, click on it with the mouse; the selection will be highlighted. Note that when you select a row, you are selecting a storage class within a particular storage subsystem.

See also the related window *Section 6.1.1, “Configured Storage Classes window”*. The *Configured Storage Classes* window lists all configured storage classes in the system, whether or not any storage resources have been assigned to them.

Field descriptions.

Active Storage Classes list

The **Column View** menu item may be used to determine which of the following columns are displayed in the list:

Subsystem

The name of the storage subsystem to which the storage class information applies. It is possible to have multiple rows in the table with identical storage class ID and Name columns, each for a different storage subsystem.

ID

The unique storage class ID number defined when the storage class is configured.

Name

The name of the storage class.

Type

The storage class type. Possible values are `Disk` or `Tape`.

Thresh

The current threshold status for the storage class. For disk, a threshold is defined as a percentage of the total bytes of storage space. When the percentage of storage space used rises above a threshold level, the threshold is exceeded. For tape, a threshold is defined as a count of free virtual volumes. When the number of free volumes drops below a threshold level, the threshold is exceeded. Possible values are:

- `OK` indicates the storage class has not reached its warning threshold.
- `Stale` indicates the storage class is active but the MPS responsible for it is shutdown, has died, or is otherwise inaccessible to SSM.

- `Warning` indicates the storage class has reached its warning threshold, but has not reached its critical threshold.
- `Critical` indicates the storage class has reached its critical threshold.

Free Space

The amount of free space left in the storage class. For disk, this is the number of free bytes. For tape, it is the number of free virtual volumes.

% Free

The **Free Space** value expressed as a percentage of **Total Space**.

Used Space

The amount of used space in the storage class. For disk, this is the number of used bytes. For tape, it is the number of used virtual volumes.

% Used

The **Used Space** value expressed as a percentage of **Total Space**.

Partially Filled VVs

The number of tapes that have been written, but have not yet reached EOM. This field will be blank for disk storage classes.

Total Space

The total amount of space in the storage class. For disk, this is the number of bytes. For tape, it is the number of virtual volumes.

Migr State

The migration state for the storage class. The **Migration State** is preserved when the MPS is recycled. The exception is `Running` which will be set to `Waiting` upon MPS recycle. Possible values are:

- `Waiting` indicates migration is not taking place at this time. The start of the next migration is waiting until criteria specified in the migration policy are met.
- `Running` indicates a migration is in progress.
- `None` indicates the storage class, as configured, is not a candidate for migration.
- `Suspended` indicates migration has been suspended.
- `Disabled` indicates migration has been stopped.

Migr Policy

The migration policy assigned to the storage class.

Purge State

The purge state for the storage class. The **Purge State** is preserved when the MPS is recycled. The exception is `Running` which will be set to `Waiting` upon MPS recycle. Possible values are:

- `Waiting` indicates purging is not taking place at this time. The start of the next purge is waiting until criteria specified in the purge policy are met.
- `Running` indicates a purge is in progress.
- `None` indicates the storage class, as configured, is not a candidate for purging.
- `Suspended` indicates purging has been suspended.
- `Disabled` indicates purging has been stopped.

Purge Policy

The purge policy assigned to the storage class.

Administration buttons.

Migration Controls

This drop-down list can be used to send migration commands to the MPS on the selected storage classes. This drop-down list is enabled only if at least one of the selected storage classes has been assigned a migration policy. Click on it to drop down a list of control options. These are:

Start

If the migration state is `Waiting`, this starts the migration and puts it into the `Running` state. This does not guarantee that any files will be migrated, only that files will be reevaluated against the appropriate migration policy and migrated as appropriate.

Stop

If the migration state is `Running`, this stops the migration and puts it into the `Disabled` state. Any ongoing I/O will be aborted.

Suspend

If the migration state is `Waiting` or `Running`, this puts the migration into the `Suspended` state. Any ongoing I/O will be aborted.

Resume

If the migration state is `Suspended`, this returns it to `Waiting` and allows MPS to again begin scheduling migration runs.

Reread policy

Tells the MPS to refresh its migration policy information by rereading the policy.

Purge Controls

This drop-down list can be used to send purge commands to the MPS on the selected storage classes. This drop-down list is enabled only if at least one of the selected storage classes has been assigned a purge policy. Click on it to drop down a list of control options. These are:

Start

If the purge state is `Waiting`, this starts the purge and puts it into the `Running` state. This does not guarantee that any files will be purged, only that files will be reevaluated against the appropriate purge policy and purged as appropriate.

Stop

If the purge state is `Running`, this stops the purge and puts it into the `Disabled` state.

Suspend

If the purge state is `Waiting` or `Running`, this puts the purge into the `Suspended` state.

Resume

If the purge state is `Suspended`, this returns it to `Waiting` and allows MPS to again begin scheduling purge runs.

Reread policy

Tells the MPS to refresh its purge policy information by rereading the policy.

Repack Volumes

Opens the *Repack Virtual Volumes* window, allowing you to start the **repack** utility program on the selected storage class. See *Section 8.4.4.1, “Repack Virtual Volumes window”* for details. This button is only enabled for tape storage classes.

Reclaim Volumes

Opens the *Reclaim Virtual Volumes* window, allowing you to start the **reclaim** utility program on the selected storage class. See *Section 8.4.4.2, “Reclaim Virtual Volumes window”* for details. This button is only enabled for tape storage classes.

Information buttons.

Storage Class Info

Opens the *MPS Storage Class Information* window for the selected storage class. For detailed information on usage of this window, see *Section 8.3.2, “MPS Disk Storage Class Information”*.

Configuration

Opens the *Storage Class Configuration* window for the selected storage class.

Migration Policy

Opens the *Migration Policy* window for the migration policy assigned to the selected storage class. If there is no assigned migration policy, the button is disabled.

Purge Policy

Opens the *Purge Policy* window for the purge policy assigned to the selected storage class. If there is no assigned purge policy, the button is disabled.

8.3.2. MPS Disk Storage Class Information

MPS Disk Storage Class Information

File Edit Help

Storage Class Name: Storage Class 1

Storage Class ID: 1

Storage Class Type: Disk

Subsystem Name: Subsystem #1

Total Space: 8,589,926,400 bytes

Free Space: 8,588,865,536 bytes

Space Used: 1 percent

Warning Threshold: 80 percent

Critical Threshold: 90 percent

Threshold Status: ● OK

Migration Attributes **Purge Attributes**

Policy Name: Purge Policy 1

Policy ID: 1

State: Waiting

Start Time: Sep 24, 2008 2:34:34 PM

End Time:

Total Units Processed: 0 bytes

Control: None ▼

Pending Operations:

☐ Freeze Update Refresh Dismiss

Retrieving data - Succeeded

This window allows you to view and update the information associated with an active disk storage class. It reports the storage space data as well as any exceeded thresholds. The window also reports detailed information on the migration and purge status. In addition, the window allows the SSM user to control the migration and purge process to override the associated migration and purge policies.

There are three differences between the disk and tape storage class information windows. First, the disk window reports space in terms of bytes while the tape window reports space in terms of virtual volumes (VVs). Second, the disk window reports *space* used while the tape window reports *percent* used. Third, the disk window shows purge status while the tape window does not.

Since an MPS is specific to a storage subsystem, the information in this window applies only to the volumes of the storage class which are assigned to a single storage subsystem. It is possible for multiple subsystems to use a storage class. If two of these windows are opened, each for the same storage class but for different subsystems, they will display different information.

MPS Disk Storage Class Information

File Edit Help

Storage Class Name: Storage Class 1

Storage Class ID: 1

Storage Class Type: Disk

Subsystem Name: Subsystem #1

Total Space: 8,589,926,400 bytes

Free Space: 8,588,865,536 bytes

Space Used: 1 percent

Warning Threshold: 80 percent

Critical Threshold: 90 percent

Threshold Status: ● OK

Migration Attributes **Purge Attributes**

Policy Name: Migration Policy 3

Policy ID: 3

State: Waiting

Start Time: Sep 24, 2008 2:34:34 PM

End Time:

Total Units Processed: 0 bytes

Control: None ▼

Pending Operations:

☐ Freeze Update Refresh Dismiss

Retrieving data - Succeeded

Any changes made to fields on the window are sent directly to the MPS and are effective immediately. If the MPS is restarted, changes made through this window are lost.

Field descriptions.

Storage Class Name

The name assigned to this storage class.

Storage Class ID

The numeric ID of this storage class.

Storage Class Type

The class of media assigned to this storage class (tape or disk).

Subsystem Name

The name of the storage subsystem for which the storage class information is being displayed.

Total Space

For disk storage classes this reports the total capacity of the storage class in bytes.

Free Space

For disk storage classes this reports the unused capacity of the storage class in bytes.

Space Used

The percent of storage space in the storage class which is used. This is calculated from **Total Space** and **Free Space** and reported as a percentage.

Warning Threshold

For disk storage classes, this value is a percentage of total storage space. When the used space in this storage class exceeds this percent of the total space, the MPS will send a warning alarm to SSM.

Critical Threshold

For disk storage classes, this value is a percentage of total storage space. When the used space in this storage class exceeds this percent of the total space, the MPS will send a critical alarm to SSM.



Note that the disk threshold reaches the warning level when the percent of used space rises above the threshold value, while the tape threshold reaches the warning level when the free volume count drops below the threshold value.

If you wish to disable the threshold so it will not generate any alarms, set this field to 100 percent for disk, and "0" volumes for tape.

Threshold Status

The current used space threshold status. Possible values are OK, warning, and Critical.

Migration Attributes tab

This panel contains the current migration status for this storage class.

Policy Name

The name of the migration policy assigned to this storage class.

Policy ID

The ID of the migration policy assigned to this storage class.

State

The current migration state for this storage class. Possible values are:

- **None** indicates this storage class, as configured, is not a candidate for migration.
- **Running** indicates a migration on this storage class is in progress.

- **Waiting** indicates migration is not taking place at this time. The start of the next migration is waiting until criteria specified in the migration policy are met.
- **Suspended** indicates migration on this storage class has been suspended by the system administrator.
- **Disabled** indicates migration on this storage class has been stopped by the system administrator.

Start Time

The date and time when the most recent migration run started. It may still be running.

End Time

The date and time when the last migration run completed.

Total Units Processed

The amount of space in the storage class which has been migrated during the current or most recent migration run. For disk storage classes and for tape storage classes running the tape file migration algorithm, this is a number of bytes.

Control

This drop-down list can be used to send commands to the MPS regarding migration on this storage class. Click on it to drop down a list of control options. These are:

None

Nothing has been selected for the migration control.

Start

If the migration state is Waiting, this starts the migration and puts it into the Running state.

Stop

If the migration state is Running, this stops the migration and puts it into the Disabled state.

Suspend

If the migration state is Waiting or Running, this puts the migration into the Suspended state.

Resume

If the migration state is Suspended, this returns it to Waiting and allows MPS to again begin scheduling migration runs.

Reread policy

Tells the MPS to refresh its migration policy information by rereading the policy.

Pending Operations

When the MPS cannot respond immediately to a control command, a command may be saved as pending. Any such pending operations are displayed here.

Purge Attributes tab

This panel contains the current purge status for this storage class. This tab appears only on the *MPS Disk Storage Class Information* window.

Policy Name

The name of the purge policy assigned to this storage class.

Policy ID

The ID of the purge policy assigned to this storage class.

State

The current purge state for this storage class. Possible values are:

- `None` indicates this storage class, as configured, is not a candidate for purging.
- `Running` indicates a purge on this storage class is in progress.
- `Waiting` indicates purging is not taking place at this time. The start of the next purge is waiting until criteria specified in the purge policy are met.
- `Suspended` indicates purging on this storage class has been suspended.
- `Disabled` indicates purging on this storage class has been stopped.

Start Time

The date and time when the most recent purge run started. It may still be running.

End Time

The date and time when the last purge run completed.

Total Units Processed

The amount of space in the storage class which has been purged, in bytes.

Control

This button can be used to send commands to the MPS regarding purging on this storage class. Click on the button to drop down a list of control options. These are:

`None`

Nothing has been selected for the purge control.

`Start`

If the purge state is `Waiting`, this starts the purge and puts it into the `Running` state.

`Stop`

If the purge state is `Running`, this stops the purge and puts it into the `Disabled` state.

`Suspend`

If the purge state is `Waiting` or `Running`, this puts the purge into the `Suspended` state.

`Resume`

If the purge state is `Suspended`, this returns it to `Waiting` and allows MPS to again begin scheduling purge runs.

`Reread policy`

Tells the MPS to refresh its purge policy information by rereading the policy.

Pending Operations

When the MPS cannot respond immediately to a control command, a command may be saved as pending. Any such pending operations are displayed here.

8.3.3. MPS Tape Storage Class Information

MPS Tape Storage Class Information

File Edit Help

Storage Class Name: Tape T10KA

Storage Class ID: 1000

Storage Class Type: Tape

Subsystem Name: Subsystem #1

Total Space: 1 VVs

Free Space: 0 VVs

Partially Filled VVs: 0

Space Used: 100 percent

Warning Threshold: 10 volumes

Critical Threshold: 5 volumes

Threshold Status: ● Critical

Migration Attributes

Policy Name:

Policy ID: 0

State: None

Start Time:

End Time:

Total Units Processed: 0

Control: None

Pending Operations:

☐ Freeze Update Refresh Dismiss

Retrieving data - Succeeded

This window allows you to view and update the information associated with an active tape storage class. It reports the storage space data as well as any exceeded thresholds. The window also reports detailed information on the migration status. In addition, the window allows the SSM user to control the migration process to override the associated migration policies.

There are three differences between the disk and tape storage class information windows. First, the disk window reports space in terms of bytes while the tape window reports space in terms of virtual

volumes (VVs). Second, the disk window reports *space* used while the tape window reports *percent* used. Third, the disk window shows purge status while the tape window does not.

Since an MPS is specific to a storage subsystem, the information in this window applies only to the volumes of the storage class which are assigned to a single storage subsystem. It is possible for multiple subsystems to use a storage class. If two of these windows are opened, each for the same storage class but for different subsystems, they will display different information.

Any changes made to fields on the window are sent directly to the MPS and are effective immediately.

Field descriptions.

Storage Class Name

The name assigned to this storage class.

Storage Class ID

The numeric ID of this storage class.

Storage Class Type

The class of media assigned to this storage class (tape or disk).

Subsystem Name

The name of the storage subsystem for which the storage class information is being displayed.

Total Space

For tape storage classes this reports the total capacity of the storage class in virtual volumes (VVs).

Free Space

For tape storage classes this reports the number of unwritten virtual volumes (VVs) in the storage class.

Partially Filled VVs

The number of tapes that have been written, but have not yet reached EOM.

Space Used

The percent of storage space in the storage class which is used. This is calculated from **Total Space** and **Free Space** and reported as a percentage.

Warning Threshold

For tape storage classes, this value is a count of free tape volumes. When the number of free volumes in this storage class drops below this level, the MPS will send a warning alarm to SSM.

Critical Threshold

For tape storage classes, this value is a count of free tape volumes. When the number of free volumes in this storage class drops below this level, the MPS will send a critical alarm to SSM.



Note that the disk threshold reaches the warning level when the percent of used space rises above the threshold value, while the tape threshold reaches the warning level when the free volume count drops below the threshold value.

If you wish to disable the threshold so it will not generate any alarms, set this field to 100 percent for disk, and "0" volumes for tape.

Threshold Status

The current used space threshold status. Possible values are OK, Warning, and Critical.

Migration Attributes

This panel contains the current migration status for this storage class.

Policy Name

The name of the migration policy assigned to this storage class.

Policy ID

The ID of the migration policy assigned to this storage class.

State

The current migration state for this storage class. Possible values are:

- `None` indicates this storage class, as configured, is not a candidate for migration.
- `Running` indicates a migration on this storage class is in progress.
- `Waiting` indicates migration is not taking place at this time. The start of the next migration is waiting until criteria specified in the migration policy are met.
- `Suspended` indicates migration on this storage class has been suspended by the system administrator.
- `Disabled` indicates migration on this storage class has been stopped by the system administrator.

Start Time

The date and time when the most recent migration run started. It may still be running.

End Time

The date and time when the last migration run completed.

Total Units Processed

The amount of space in the storage class which has been migrated during the current or most recent migration run.

Control

This drop-down list can be used to send commands to the MPS regarding migration on this storage class. Click on it to drop down a list of control options. These are:

None

Nothing has been selected for the migration control.

Start

If the migration state is `Waiting`, this starts the migration and puts it into the `Running` state.

Stop

If the migration state is `Running`, this stops the migration and puts it into the `Disabled` state.

Suspend

If the migration state is `Waiting` or `Running`, this puts the migration into the `Suspended` state.

Resume

If the migration state is Suspended, this returns it to Waiting and allows MPS to again begin scheduling migration runs.

Reread policy

Tells the MPS to refresh its migration policy information by rereading the policy.

Pending Operations

When the MPS cannot respond immediately to a control command, a command may be saved as pending. Any such pending operations are displayed here.

8.4. Dealing with a space shortage

If free space warning thresholds are exceeded in a storage class, HPSS will warn the administrator with messages in the *Alarms and Events* window. Before the available storage space in a storage class drops to a level where the storage class may no longer be able to satisfy requests, efforts should be made to create additional free space. This can be accomplished by three methods.

The first method is to use the migration and purge functions to free up space. Migration and purge can be run manually; see *Section 8.4.1, “Manually starting migration”* and *Section 8.4.2, “Manually starting purge”*. For a long term solution, it may be desirable to tune the migration and purge policies to cause the automatic migration and purge operations to occur more frequently. This method of increasing storage space works for all levels in the Class of Service hierarchy except the lowest level.

The second method is to add additional storage volumes to the storage class using the normal import and create resources procedures. See *Section 8.1, “Adding storage space”*.

The third method is to increase free tape storage space by repacking sparse volumes and reclaiming empty volumes. The **repack** program moves information from sparsely filled volumes to other volumes in the same storage class. When the **repack** is finished, a number of volumes will be cleared of data. To make these volumes ready for reuse, they must be processed by the **reclaim** program, which refreshes the Core Server metadata that describes the volumes. See *Section 8.4.3, “Repacking and reclaiming volumes”*.

8.4.1. Manually starting migration

The Migration/Purge Server runs migration periodically in the time interval specified in the migration policy. However, between these automatic migration runs, an administrator can use the *Active Storage Classes* window to start a migration manually. When a migration run is started manually, the run timer is reset. The next automatic migration will take place after the specified interval after the end of the manual run.

From the *Active Storage Classes* list window, select one or more migratable storage classes and then click on the **Migration Controls** drop-down button and select **Start** from the list.

Note: This request will go straight to the Migration/Purge Server without confirmation.

The Migration/Purge Server stops the migration run when either the Migration Target threshold in the migration policy is reached or there are no more disk files or tape virtual volumes eligible for migration.

The only difference between a manually-started migration run and an automated migration run is that the manually-started run is initiated manually by an administrator. Therefore, a manually-started migration will not migrate files that do not meet the migration criteria.

8.4.2. Manually starting purge

The Migration/Purge Server obtains the disk storage class statistics periodically. The period is set in the Migration/Purge Server's specific configuration record. The Migration/Purge Server then evaluates the used space and may start a purge run. Between these automatic purge runs, an administrator can use the *Active Storage Classes* window to start the purge run manually for one or more storage classes.

From the *Active Storage Classes* window, select one or more purgeable storage classes, and then click on the **Purge Controls** drop-down button and select **Start** from the list.

Note: This request will require confirmation before going to the Migration/Purge Server.

The Migration/Purge Server stops purging when the space used falls to the value specified in the purge policy or there are no more disk files eligible for purging.

The only difference between a manually-started purge run and an automated purge run is that the manually-started run is initiated manually by an administrator. Therefore, a manually-started purge will stop purging when the appropriate threshold is reached just as an automated purge run would.

8.4.3. Repacking and reclaiming volumes

Over time, the active data on a tape virtual volume may become sparse as files are deleted or replaced. The **repack** utility program provides the administrator with the ability to move data from sparse volumes to other volumes, allowing more efficient use of tape resources. **Repack** is frequently used to migrate files from older tapes to newer tapes when new media replaces older media. The program can be invoked via the command line or from an SSM window. Running **repack** from the command line affords the greatest flexibility in its use and the opportunity to monitor its progress.

Repack can also move disk data from one volume to another. **Repack** is frequently used to empty a failing disk volume. Begin by setting the volume to RW (or perhaps RO depending on the nature of the failure) using the SSM interface, then move its contents to other volumes using **repack**.

Repack selects tape volumes in one of two ways. The administrator can provide a list of tape volumes to repack, or **repack** can select volumes based on a number of selection criteria. If **repack** is provided with a list of tape volumes to process, those volumes must be in RW, RO, EOM or EMPTY condition. Volumes in RWC or DOWN condition cannot be selected. Repacking RW and EMPTY tape volumes is permitted because the administrator has explicitly selected the tape. Repacking EMPTY tape volumes is also permitted because it is harmless to do so and it provides a way to prove that the volume is empty.

When **repack** selects tape volumes to process, it uses *number*, *storage class*, *family*, *retired*, *checked out*, and *percentage full* criteria to select a set of tape volumes. The selected volumes will be in RO or EOM condition. Volumes in RWC, EMPTY or DOWN condition are never selected. Volumes that are in RW condition and are also Retired will be treated as if they are in RO condition and may be selected if they qualify under the other selection criteria. If the *--select-retired-tapes* option is used, only retired tapes that meet the rest of the selection criteria will be selected.

When repacking tapes, by default **repack** will only copy active data to repack output tapes. This means it will skip over inactive data (also called whitespace) on source tapes.

Repack also creates new aggregates on the output tape by bringing together smaller tape storage segments from the input tape and writing new aggregates in the output tape. The `--max-bytes-per-aggregate` command line option controls how this is done. Aggregation can be disabled by specifying `--max-bytes-per-aggregate=0` on the command line. The storage class VV *Estimated Size* is the default value and upper limit for this command line option.

The default behavior of creating new aggregates on the output tape and removing whitespace can be turned off via the `--reduce-tape-positioning` or `--reduce-tape-pos-detect` options. It has been observed that reorganizing the data within a sparse tape (a tape with a significant amount of whitespace) may incur a heavy performance penalty (up to 2-3x) due to extra tape positioning. In certain cases, it may be beneficial to skip this reorganization process in order to achieve faster repack rates. However, it is worth noting that eventually a repack without `--reduce-tape-positioning` should be done to reclaim the whitespace being carried around within the aggregates.

When repacking RAIT volumes, the program accepts the `--select-degraded-rait` command line option. When this option is provided, **repack** selects tape volumes that contain one or more physical volumes that have had mount, read or write errors. Specific details regarding such errors are recorded in metadata and are accessible via the **lspvhist** utility program. This program can help to shed further light on whether an apparently degraded RAIT volume should be repacked.

Tape **repack** can segregate repack output tapes from other tapes. When this feature is enabled, **repack** will copy files only to tapes that contain repack output. New files migrated to tape will be written on other tapes. This feature is controlled by the `HPSS_CORE_REPACK_OUTPUT` environment variable. If this variable is present in the system's environment, and if its value is "on", tape repack output tapes will be segregated from other tapes. This feature is "on" by default. Set `HPSS_CORE_REPACK_OUTPUT` to "off" to disable it system-wide. As an alternative to disabling the feature system-wide, this feature can be turned off for a specific **repack** instance. This is controlled by **repack**'s `--no-dest-tape-segregation` command line option. See the **repack** man page for more details.

Disk volumes to be repacked must be in RW, RO or EMPTY condition. EMPTY volumes may be repacked because it is harmless to do so, the administrator has explicitly selected the disks, and it provides a way to prove that the volume is empty.

Repacking disks cannot be done via the SSM window; the program must be executed from a command line following the instructions given in the **repack** man page.

When the **repack** process completes successfully, the Number of Active Segments in the volume's display will be zero and the VV Condition will be EMPTY.

Note that tape volumes in EOM condition go to EMPTY condition when **repack** moves all of the data on the volumes to other volumes. Volumes in RW and RO condition do not go to EMPTY condition. Manually setting the condition of an empty tape volume to EOM causes it to go to EOM condition, then immediately to EMPTY condition.

After a tape volume is successfully repacked, it remains in EMPTY state and is not available for storing data. To make the volume available, it must be reclaimed. Disk volumes can be placed back in service by changing the VV Condition to RWC.

Note that both disk and tape repack operations can put significant demands on system resources, particularly when repacking striped volumes. Since **repack** performs a tape-to-tape copy operation, it will consume twice as many tape drives as the stripe width of the volumes being repacked.

The **reclaim** program deletes empty tape virtual volumes and re-creates them with the same configuration as the original volume. The PV Name(s), Stripe Width, PV and VV Block Sizes, PV and VV Estimated Size and other parameters of the old VV are used to create a new VV with a new Virtual Volume ID (VVID). The metadata describing the old VV is deleted. HPSS requires that tapes be reclaimed, rather than reused, so that it can be certain that no references to the old VV exist or could be resolved.

Reclaimed tape volumes are entered into HPSS metadata in RWC condition and are immediately available for use.

8.4.4. Verifying data integrity

An administrator may want to verify data which has been sitting at rest on a tape volume periodically or in response to some event. This can be accomplished by using the **lbp_verify** tool.

lbp_verify causes a tape volume to be mounted, and segments read and verified against a CRC stored on the medium. The tool provides a number of options for specifying inputs and drive resources to use. See the `lbp_verify.7` man page for more details.

While the tool is verifying the volume, no other I/O will be processed against the volume. When the verification is complete, an entry will be logged in the PV history table and if all PVs in the VV were successfully verified, the volume's last verify time will be updated. If an error is found, information about the location of the error will come back from the tool and be logged in the PV history table.

Verification uses a checkpoint system to keep track of progress against a tape. If for some reason verification cannot be completed, the next time the volume is requested to be verified, the process can be started again from the checkpoint.

8.4.4.1. Repack Virtual Volumes window

Repack Virtual Volumes

File Edit Help

Storage Class Name

Storage Class ID

Subsystem ID

Core Server

File Family Criteria

Number of VVs to Repack

VV Space percent

Max Bytes in Aggregate (0 to disable file aggregation)

Estimated Size

Repack Options

☐ Select Only Retired VVs

☐ Include Shelved VVs

This window provides an interface for repacking tape virtual volumes.

Field descriptions.

Storage Class Name

The name of the storage class that will be repacked.

Storage Class ID

The ID of the storage class that will be repacked.

Subsystem ID

The ID of the subsystem which contains the volumes to be repacked.

Core Server

The name of the Core Server that manages the volumes to be repacked.

File Family Criteria

Indication of whether to use file family criteria. Click on the button to drop-down a list of available options. Select `Ignore` to indicate not to use the file family criteria when selecting

volumes for repack. Select the name of a file family to repack if you want to use file family criteria. When a file family is selected, volumes assigned to that family will become candidates for repack. Volumes not in the selected family will not be candidates. Storage segments on repacked volumes will be transferred to other volumes assigned to the same family. If no writable tape volumes are available in the family, a new blank volume will be chosen and assigned to the family. If `Not in a family` is selected, volumes not assigned to any file family will become candidates for repack. Volumes that have been assigned to a file family will not be candidates.

Number of VVs to Repack

The maximum number of tape volumes to repack. If **repack** does not find this many candidate volumes, those that have been selected will be repacked and no error will result.

VV Space

This is a repack selection criterion expressed as a percentage. If the amount of occupied space on a candidate tape is less than this value, the tape will be considered for repacking. Tape volumes with larger amounts of occupied space will not be considered.

The occupied space percentage is calculated by summing the lengths of the files on the candidate tape and comparing that sum to the VV Estimated Size field in the storage map. If this ratio is less than the **VV Space** percentage, the volume is selected. Note that vacated files in file aggregates on tapes will not be counted as occupied space. A tape containing a relatively small number of large file aggregates from which a large percentage of the files have been deleted, may be selected by **repack**.

If **VV Space** is 100, the comparison is not performed and selection of tape volumes to repack is made using the remaining criteria.

Max Bytes in Aggregate

This sets a limit on the number of bytes that will be accumulated into a tape aggregate written by **repack**. This value is limited to the configured **Estimated Size** of VV's in the selected storage class. Set this value to zero to disable the creation of new aggregates in output tapes. This value has no effect when repacking disks.

Estimated Size

The estimated size of the storage class to be repacked. This value comes from the storage class configuration and is for informational purposes only to help in setting a valid **Max Bytes in Aggregate** value.

Repack options

Select Only Retired VVs

If selected, **repack** selects only retired volumes in the indicated storage class. Retired volumes may be in RO, EMPTY or EOM condition. DOWN volumes are never selected. If this option is not selected, both retired and non-retired tape volumes in RO, EMPTY or EOM condition may be selected.

Include Shelved VVs

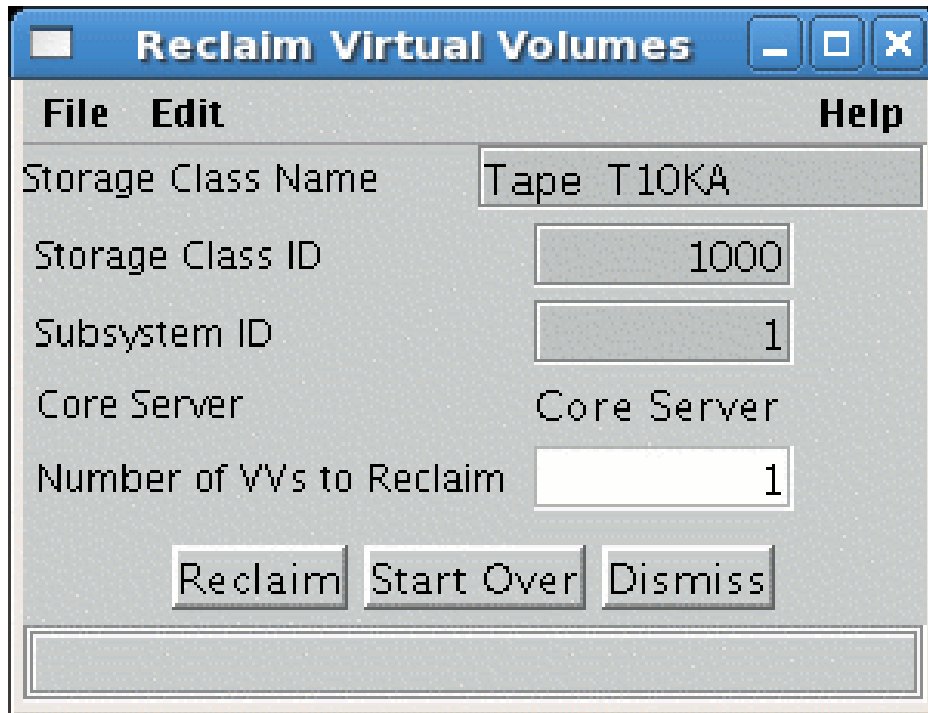
Tape volumes which have been removed from HPSS using the shelf tape utility are eligible for repack.

Buttons.

Repack

Click this button to start the **repack** program on the indicated storage class. Tape volumes that meet the selection criteria will be repacked. Status messages are displayed on the status bar at the start and end of the repack.

8.4.4.2. Reclaim Virtual Volumes window



This window is used to run the **reclaim** utility program which removes and recreates tape virtual volumes.

Field descriptions.**Storage Class Name**

The name of the storage class selected for reclaiming.

Storage Class ID

The ID of the storage class selected for reclaiming.

Subsystem ID

The ID of the subsystem in which the volumes to be reclaimed reside.

Core Server

The name of the Core Server that will perform the reclaim operations.

Number of VVs to Reclaim

The number of tape virtual volumes in the storage class that will be reclaimed. If the number of candidate volumes is less than this value, all eligible volumes will be reclaimed. Retired volumes are never reclaimed.

Buttons.

Reclaim

Click this button to start the **reclaim** utility program on the indicated storage class. Tape volumes that are described as EMPTY, and are not retired, will be reclaimed. Status messages are displayed on the status bar at the bottom of the window at the start and end of the reclaim.

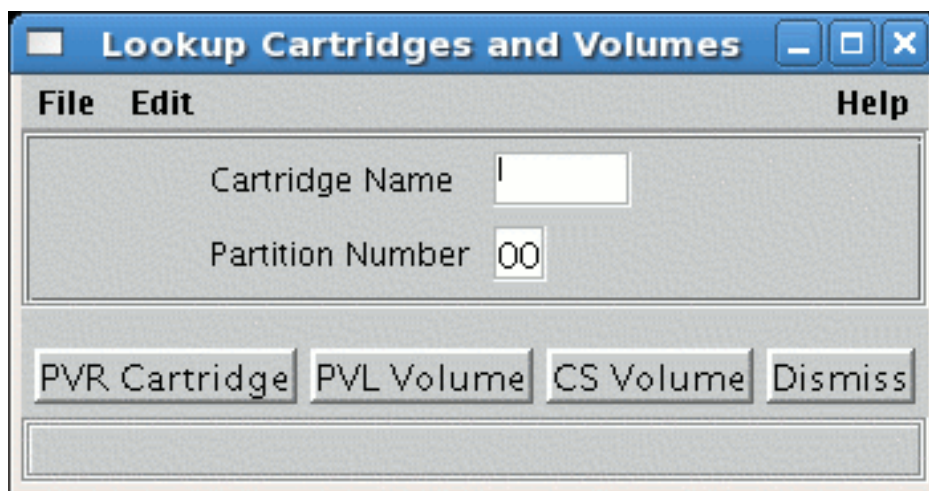
SSM invokes the **reclaim** utility program and passes it the storage class and number of tape volumes to be reclaimed, as entered on the *Reclaim Virtual Volumes* window. The utility program selects empty tape volumes from the storage class until it selects the required number or the supply of empty volumes is exhausted. The volumes are reclaimed and immediately become available for use.

The **reclaim** utility can also be invoked as a command line utility. Used this way, the user has more flexibility in choosing options and can monitor the program's progress. Refer to the **reclaim** man page for more information on how to invoke the utility.

8.5. Volume management

This section describes how to monitor and manage PVL volumes, PVR cartridges, and Core Server disk and tape volumes.

8.5.1. Lookup Cartridges and Volumes window



This window allows you to access information windows for PVR cartridges, PVL volumes, and Core Server volumes. After clicking a button to perform a lookup operation, this window will remain open to allow additional lookups.

Field descriptions.**Cartridge Name**

The six-character cartridge label. A string exactly six characters long must be entered in this field.

Partition Number

This field is no longer used and should always be "00".

Buttons.

PVR Cartridge (tape only)

Once you have filled in the cartridge name, clicking on this button will open the *PVR Cartridge Information* window for the specified cartridge. This metadata is created when the cartridge is successfully imported.

PVL Volume

Once you have filled in both fields, clicking on this button will open the *PVL Volume Information* window for the specified volume. This metadata is created when the volume is successfully imported.

CS Volume

Once you have filled in both fields, clicking on this button will open the *Core Server Disk Volume* or *Core Server Tape Volume* window for the specified volume. This metadata is created when the disk or tape storage resources are successfully created.

8.5.2. PVL Volume Information window

The *PVL Volume Information* window allows the SSM user to view the data for imported volumes.

Before using the window, the user should know the six-character label of the PVL volume. The **dumpppv_pvl** utility can be used to list all the volumes being managed by the PVL and help determine the label.

The screenshot shows the 'PVL Volume Information' window with the following fields and values:

Field	Value
Volume ID	HPSS0600
PVR Server	
Allocated Client	Core Server
Allocation State	Allocated
Volume Type	Generic - Default Disk
Label Format	HPSS
Operational State	Enabled (indicated by a green dot)
Usage State	Busy (indicated by a green dot)
Administrative State	Unlocked
Cartridge ID	HPSS06
Comment	

At the bottom of the window, there are four buttons: ☐ Freeze, Update, Refresh, and Dismiss. Below these buttons is a status bar that reads 'Retrieving data - Succeeded'.

This window allows you to view the information associated with a PVL volume.

Field descriptions.**Volume ID**

The eight-character volume label for the volume being viewed.

PVR Server

The descriptive name of the PVR that owns the cartridge corresponding to the volume being viewed. This field is applicable only to tape volumes and is always blank for disk volumes.

Allocated Client

The descriptive name of the Core Server which manages the volume. This is the name of the Core Server specified when storage resources were allocated on the volume using the Create Disk Resources or Create Tape Resources operation. If no storage resources have been created on the volume, this field will be blank.

Allocation State

The usability of the volume. This includes whether the volume import and label operation completed successfully, whether storage resources have been created on the volume, and whether the volume is currently stored in the robot or removed to an external shelf or vault.

Possible values are:

- **Allocated** indicates storage resources have been created on the volume, it is allocated to a Core Server, and it is available to the system for I/O.
- **Unallocated** indicates the volume has been successfully imported into the HPSS system and labeled. However, no storage resources have been created on it and it has not been allocated to a Core Server. It is therefore not available to the system for I/O.
- **Allocated - On Shelf** indicates storage resources have been created on the volume, it is assigned to a Core Server, and it is available to the system. However, it has been removed from the tape robot (to a vault, for example), so there will be a delay for any mount operation while the volume is returned to the robot.
- **Unlabeled** indicates an error occurred during the mount/label operation during import. A volume in this state is not usable by HPSS. It should be exported and re-imported.
- **Allocated - No Shelf Info** indicates the volume has been allocated to a Core Server and is available to the system for I/O, but no shelf information is available because the PVL can not communicate with the PVR.

Volume Type

The HPSS media type corresponding to this volume.

Label Format

The type of label on the volume. Possible values are:

- **HPSS** indicates the volume has an ANSI label (it starts with an 80-byte block beginning with the characters "VOL1"). The owner field of the label contains "HPSS".
- **Foreign** indicates the volume has an ANSI label, but the owner field is something other than "HPSS".
- **None** indicates the volume has no label. A volume in this state is not usable by HPSS. It should be exported and re-imported.

- `Data` indicates the volume is an uninitialized tape.
- `NonANSI` indicates the label is an 80-byte non-ANSI record.
- `NSL` indicates the volume has a NSL label format (80-byte record without EOF).
- `Unknown` indicates the label of the volume failed. Status of label unknown.

Operational State

This will always be `Enabled`.

Usage State

`Busy` if the volume is in use; `Idle` otherwise.

Administrative State

This will always be `Unlocked`.

Cartridge ID

The six-character cartridge label for the volume.

Comment

This field provides a 128-character buffer in the PVL volume metadata which gives the administrator the opportunity to associate miscellaneous text with a volume. For example, a site may want to place a comment in this field that the volume (cartridge or disk) is out of service.

8.5.3. PVR Cartridge Information window

The *PVR Cartridge Information* window allows the SSM user to view the data about a cartridge managed by one of the HPSS PVRs.

Before using the window, the user should know the six-character label of the cartridge. The **`dumppv_pvr`** utility can be used to list all the cartridges being managed by the PVR and help determine the label.

PVR Cartridge Information

File Edit Help

Cartridge ID: X02030

Sides: 1

Mounted Side: 0

PVR Server: Operator PVR

Cartridge Type: STK - T10000A Tape

Manufacturer:

Lot Number:

Service Start Date: May 15, 2008 10:10:56 AM

Maintenance Date: Oct 30, 2008 8:00:00 AM (ex. Jan 31, 2009 10:59:30 AM)

Last Mounted Date: May 15, 2008 11:25:23 AM

Mounts In Service: 1

Mounts Since Maintenance: 0

Operational State: ☒ Enabled

Usage State: ☒ Active

Administrative State: Unlocked

Mount Status: ☒ Dismounted

Job Id: 0

Location

Type: None

Port: 0

Drive: 0

Slot

Unit: 0 Panel: 0 Row: 0 Column: 0

☐ Freeze

Retrieving data - Succeeded

This window allows you to view and update the information associated with an HPSS tape cartridge.

Note that the **Location Type** fields are represented differently for certain types of robots, for which **Port**, **Drive** and **Slot (Unit, Panel, Row, and Column)** may each be displayed as "0". If it is necessary to locate a cartridge in one of these robots, the robot's operator interface must be used.

This window contains three bookkeeping fields: **Maintenance Date**, **Mounts Since Maintenance**, and **Mount Status**. If the some maintenance is performed on the cartridge, such as re-tensioning it,

the **Mounts Since Maintenance** field can be reset. The **Maintenance Date** field can also be reset. Note that resetting one field does not affect the other. The third operation involves resetting the **Mount Status** from Mount Pending to Dismounted via the **Cancel Mount** button. This operation is only required for handling a specific error condition.

Any changes made to fields on this window are sent directly to the PVR and are effective immediately.

Field descriptions.

Cartridge ID

The six-character cartridge label.

Sides

The number of partitions (sides) on the cartridge. Currently always set to "1".

Mounted Side

The partition currently mounted. Set to zero if the cartridge is not currently mounted; note that this can also mean that side 0 of the cartridge is mounted.

PVR Server

The descriptive name of the PVR which manages the cartridge.

Cartridge Type

The HPSS media type corresponding to this cartridge. This controls which the type of drive in which the cartridge can be mounted.

Manufacturer

The Manufacturer string specified when the cartridge was imported.

Lot Number

The Lot Number string specified when the cartridge was imported.

Service Start Date

The date and time when the cartridge was imported.

Maintenance Date

If the maintenance is performed on the cartridge, for example by retensioning it, this field can be set to record the date of the maintenance. The format for entering the date is given to the right of the field. This format can be changed with the **-D** option to the **hpssgui** startup script.

Last Mounted Date

The date and time when the cartridge was last mounted.

Mounts In Service

The number of times the cartridge has been mounted by HPSS since it was imported.

Mounts Since Maintenance

The number of times the cartridge has been mounted since its last maintenance. Click the **Reset** button to the right of this field to reset the value to zero. This field is not automatically reset when Maintenance Date is changed.

Operational State

This will always be `Enabled`.

Usage State

This should always be `Active`.

Administrative State

This will always be `Unlocked`.

Mount Status

The mount status of the cartridge. Possible values are:

- `Mount Pending`
- `Mounted`
- `Dismount Pending`
- `Dismounted`
- `Eject Pending`
- `Checked In`
- `Check-In Pending`
- `On Shelf`
- `Check-Out Pending`
- `Move Pending`
- `Move Shelf Pending`

Job Id

The job identifier associated with the cartridge. A value of zero indicates that there is not a job associated with the cartridge (that is, when the cartridge is not mounted).

Location Type

The type of cartridge location information being displayed. Possible values are:

- `None` indicates no location information is available.
- `Port` indicates the location is a port number. [This option is currently not used.]
- `Drive` indicates the location is a drive ID number.
- `Slot` indicates the location is a slot specification.

The following fields are filled in (nonzero) based on the **Location Type** and whether or not the PVR has the information:

Port

The port number where the cartridge is located. This option is currently not used and thus will always be zero.

Drive ID

The drive ID number where the cartridge is located. This field is valid when the **Location Type** is Drive. It is used/valid when the cartridge is Mounted; it is not used/valid (and thus will have the value of zero) when the cartridge is not Mounted.

Slot

The slot address (**Unit, Panel, Row** and **Column**) for the cartridge. These fields are valid only if **Location Type** is Slot and the PVR supports this (STK and AML). The fields will be zero when it's not valid.

Buttons.

Cancel Mount

When **Mount Status** is displaying Mount Pending, the **Cancel Mount** button to the right of the field will be sensitized. If you click the button, a large warning will be displayed, and you will have to confirm your request by clicking **Cancel Mount** again. This will send a request to the PVR to cancel the mount request for the cartridge. As noted in the on-screen warning, this should only be done as a last-ditch attempt to cancel a mount request which cannot be canceled in any other way.



Note that this will not dismount a mounted tape.

8.5.4. Core Server Volume and Segment windows

8.5.4.1. Core Server Disk Volume information window

The screenshot shows the 'Core Server Disk Volume' window. At the top, the title bar says 'Core Server Disk Volume'. Below it are 'File', 'Edit', and 'Help' menus. The main area contains the following fields:

- Name:** DSK10300
- VV Condition:** RWC (dropdown menu)
- Retired:** ☐
- Changes Pending:** None
- Offline:** ☐

Below these are two tabs: 'More commonly used data' (selected) and 'Less commonly used data'. The 'More commonly used data' tab contains:

- Number of Extents:** 0
- Storage Class:** 105 | 4wd (classic)
- Storage Map Flags:** (empty field)
- Usable Length:** 65,305,313,280
- Free Space:** 65,305,313,280
- Stripe Width:** 4
- Time Last Read:** (empty field)
- Time Last Written:** (empty field)
- PVL Job ID:** 2

Below the tabs is a section titled 'Physical Volumes' containing a table:

Vol Name ▼	Type	Dev ID	Mvr IP Addr
DSK10300	Generic - Default Disk	209	0.0.0.0 [0]
DSK20300	Generic - Default Disk	210	0.0.0.0 [0]
DSK30300	Generic - Default Disk	211	0.0.0.0 [0]
DSK40300	Generic - Default Disk	212	0.0.0.0 [0]

At the bottom of the window are buttons: ☐ Freeze, Update, Refresh, and Dismiss. A status bar at the very bottom says 'Retrieving data - Succeeded'.

This window displays information about a disk volume as represented by the Core Server.

Field descriptions.

Name

The ASCII name of the first physical volume that is a part of the disk virtual volume. The entire virtual volume can be referred to by this name.

VV Condition

This is the administrative control for the disk virtual volume. It will have one of five values: RWC, RW, RO, EMPTY or DOWN.

- In RWC condition, the volume can be read and written. This is the normal operational state.
- In RW condition, the volume can be read and written, but new disk storage segments may not be created on the volume.
- In RO condition, the volume can be read but not written. New storage segments cannot be created on the volume.
- In EMPTY condition, the volume cannot be read or written. All storage segments have been removed from the volume, and new segments cannot be created. Since this condition represents a physical state of the volume, the operator may not change the volume to this condition. EMPTY is displayed by the server when the condition is RO and there are no storage segments in the volume.
- In DOWN condition, the volume cannot be read, written or mounted. This condition can be used to make a disk unavailable to the system.

Change the **VV Condition** of a disk virtual volume by selecting the desired condition from the drop-down list and then clicking the **Update** button.

Changes Pending

If there are any VV Condition changes for this volume pending in the Core server, **Changes Pending** will be `Yes` with a red bullet. Otherwise, **Changes Pending** will be `None`.

Retired

When checked, the volume is retired. New storage segments will not be created on the volume.

Offline

When checked, the volume is considered offline and not usable for reading nor writing. This often occurs as a result of failed communication with the Mover that controls the volume. HPSS will attempt to automatically return offline disks to service. For disks that HPSS can automatically set back online, you will see corresponding entries in the **hpssgui Alarms and Events** window. Because this occurs in the background, failures to automatically set disk volumes back online are silent.

If a disk is functioning sporadically, the disk may transition back and forth between offline and online. If you want to take the disk out of service for repairs, set its VV Condition to DOWN.

A reminder about offline disks will appear in the **hpssgui Alarms and Events** window on a regular basis. By default, the reminder message appears once a day in the mid-morning. You can alter the frequency with which the message appears by setting the `HPSS_OFFLINE_DISK_REMINDER_INTERVAL_HOURS` environment variable. Note that the value you specify should be a number of hours, not minutes nor seconds.



Offline versus DOWN

Disks that are either offline or have a VV Condition of DOWN are unusable for read and write activity. They are essentially *out of service*. Though both states, offline and DOWN, achieve the same end result, they reflect different situations.

The offline state reflects a physical reality (for example, the disk's Mover may be down or not reachable over the network). It is automatically managed by HPSS and is

not meant to be used to administratively control access to disk volumes. In general, if you want a disk to be unusable, *set its VV Condition to DOWN rather than setting the Offline state.*

DOWN is an administrative state. It's useful for removing a disk from service regardless of whether it has an apparent malfunction.

More commonly used data tab

Number of Extents

The number of disk storage segment extents allocated in this virtual volume.

Storage Class

The storage class to which the disk virtual volume is assigned.

Storage Map Flags

This field is usually blank for volumes that are actively being used to store incoming data. Volumes in other states (such as, but not limited to, full tapes, tapes with no data yet written to them, tapes written to by repack, offline disks) may contain a status in the field. This field is informational only.

Usable Length

The amount of disk space that can be used for storing storage segments. This is the Actual Length of the disk volume minus certain amounts set aside for system use.

Free Space

The number of bytes on the disk virtual volume that are not assigned to storage segments. **Free Space** is initially set to the value of **Usable Length**, then decremented and incremented by the size of disk storage segments as they are created and deleted. Since disk storage segments are allocated in whole clusters, they are usually larger than the amount of storage needed to store the segment of the file. This creates "slack space" which is space allocated to files, but not filled with file data. The amount of disk space in use for any volume, which is the **Usable Length** minus the **Free Space**, will always be somewhat greater than the sum of the lengths of the file segments stored on the volume. Since unallocated disk space may become fragmented, it may not be possible to create new disk storage segments on a volume even if the amount of free space appears to be large. The size of the largest free extent on the volume is logged by the Core Server when this happens.

Stripe Width

The number of physical volumes in the stripe group that makes up the virtual volume. All volumes in HPSS are considered to be striped, even if this value is one.

Time Last Read

The date and time the volume was last read. This value is not stored in metadata, but is initialized to zero when the Core Server starts and then updated each time the disk is read. This field is not affected by transaction aborts.

Time Last Written

The date and time the volume was last written. This value is not stored in metadata, but is initialized to zero when the Core Server starts and then updated each time the disk is written. This field is not affected by transaction aborts.

PVL Job ID

The PVL job ID in which this volume is mounted. This field will be zero if the volume is not mounted.

Core Server Disk Volume

File Edit Help

Name:

VV Condition: ☐ Retired

Changes Pending: ☐ Offline

More commonly used data **Less commonly used data**

Actual Length:

PV Size:

Cluster Length:

VV BlockSize:

PV BlockSize:

Creation Time:

VV ID:

Object Id:

Physical Volumes

Vol Name ▼	Type	Dev ID	Mvr IP Addr	Mvr	Mvr
D0123300	Generic - Default Disk	2233	0.0.0.0 [0]	Mover (hpss-mvr20)	
D0123400	Generic - Default Disk	2234	0.0.0.0 [0]	Mover (hpss-mvr20)	
D0123500	Generic - Default Disk	2235	0.0.0.0 [0]	Mover (hpss-mvr20)	
D0123600	Generic - Default Disk	2236	0.0.0.0 [0]	Mover (hpss-mvr20)	

4/4

☐ Freeze

Retrieving data - Succeeded

Less commonly used data tab**Actual Length**

The length of disk virtual volume in bytes. This length includes all of the space set aside for system use. See **Usable Length**.

PV Size

The length in bytes of a physical volume in this disk virtual volume. All physical volumes in a disk virtual volume must be the same length.

Cluster Length

The size of the allocation unit, in bytes. When disk storage segments are created on the volume, they are created at a length that will be a multiple of this value.

VV Block Size

The virtual volume block size. This is the number of bytes written from a data stream to an element of the striped volume before the stream switches to the next element of the stripe.

PV Block Size

The size, in bytes, of the media data block.

Creation Time

The date and time the disk volume was created in the Core Server.

VV ID

The detailed description of the disk virtual volume ID associated with this disk virtual volume.

Physical Volumes

This is a table of physical volume attributes for the physical volumes that make up this disk virtual volume.

Vol Name

The ASCII name of the physical volume.

Type

The media type.

Dev ID

The ID of the device the physical volume is mounted on.

Mvr IP Addr

The IP address of the Mover that operates this physical volume.

Mvr

The descriptive name of the Mover that operates this physical volume.

Mvr Host

The name of the host on which the Mover runs.

8.5.4.2. Core Server Tape Volume information window

The screenshot shows the 'Core Server Tape Volume' window. At the top, the title bar says 'Core Server Tape Volume'. Below it are 'File', 'Edit', and 'Help' menus. The main area contains fields for 'Name' (37312000), 'VV Condition' (RWC), 'Retired' (checkbox), 'Changes Pending' (None), and 'FIFO' (checkbox). There are three tabs: 'More commonly used data' (selected), 'Less commonly used data', and 'RAIT Options'. Under 'More commonly used data', there are fields for 'Number of Active Segments' (1), 'Map State' (Free), 'Storage Class' (2 | 1wt (tape)), 'Map Flags', 'Max Written Length' (1,743,781,888), 'Active Length' (1,743,781,888), 'Estimated Size' (1,717,986,918,400), 'Space Left' (1,716,243,136,512), 'Number of Reads' (0), 'Number of Writes' (4), 'Time Last Read', 'Time Last Written' (Jan 29, 2016 8:03:18 PM), 'Time Last Verified' (Feb 1, 2016 5:05:56 PM), and 'PVL Job ID' (0). Below these are 'Current Position' fields for 'Section' (0) and 'Offset' (0), and 'Next Write Address' fields for 'Section' (1) and 'Offset' (1,743,781,888). At the bottom, there is a 'Physical Volumes' table with columns: Vol Name, Type, Dev ID, Condition, Errors, and Format. The table has one row: 37312000, IBM - 3592 E07 Standard Tape, 0, RW, and HI. Below the table are buttons for 'Freeze', 'Update', 'Refresh', and 'Dismiss'. At the very bottom, a status bar says 'Retrieving data - Succeeded'.

Vol Name	Type	Dev ID	Condition	Errors	Format
37312000	IBM - 3592 E07 Standard Tape	0	RW		HI

This window displays information about a tape volume as represented by the Core Server.

Field descriptions.

Name

The ASCII name of the first physical volume that is a part of the tape virtual volume. The entire virtual volume can be referred to by this name.

VV Condition

This is the administrative control for the tape virtual volume. It will have one of five values: RWC, RO, EOM, EMPTY or DOWN.

- In RWC condition, the volume can be read and written. This is the normal operational state.
- In RO condition, the volume can be read but not written. New storage segments cannot be created on the volume.
- In EOM condition, the volume can be read but not written. One or more of the tapes has been written to its end and the tape virtual volume is now full. New storage segments cannot be

created on the volume. The volume condition can be changed to DOWN, but not to any other condition.



Note that setting a volume to EOM is irreversible. Don't use this setting unless you're sure it's what you want.

- In EMPTY condition, the volume cannot be read or written. The volume has reached EOM, and all storage segments have been removed from the volume. New storage segments cannot be created. Volumes cannot be changed manually to EMPTY condition. EMPTY is reported by the Core Server when the volume has reached EOM and no storage segments remain on the volume.
- In DOWN condition, the volume cannot be read, written or mounted. This condition can be used to make a tape unavailable to the system.

Change the condition of a tape virtual volume by selecting the desired condition from the drop-down list and then clicking the **Update** button.

Changes Pending

If there are any VV Condition changes for this volume pending in the Core server, **Changes Pending** will be Yes with a red bullet. Otherwise, **Changes Pending** will be None.

Retired

When checked, the volume is retired. New storage segments will not be created on the volume. Retired volumes may not be reclaimed.

FIFO

When checked, the Core Server will schedule read requests on the volume in FIFO order, regardless of the Core Server's tape queue method settings. The VV's queue will not become FIFO-ordered until the next time the queue is ordered — generally due to a new request coming in.

More commonly used data tab

Number of Active Segments

The number of storage segments stored in this virtual volume.

Map State

This is the primary operational state flag of the tape virtual volume. It can take on one of the following values:

- Free indicates the volume is available to be written.
- Allocated indicates the volume is assigned to a tape write operation and is being written.
- EOM indicates the volume has reached End Of Media and can no longer be written.
- Empty indicates the volume has reached End Of Media and all storage segments have been deleted.
- Deny indicates the volume is not eligible to have a new tape storage segment created on it. Tape read operations are permitted, and the last storage segment on the tape may be written under certain circumstances.

The **Map State** field is displayed to provide an additional level of detailed information about the tape volume. It cannot be changed manually. Use the **VV Condition** to change the operational condition of the volume.

A tape volume starts in `Free` state, which indicates that it can be selected to receive a storage segment and its data. When data is being written to a volume, it will be shown in `Allocated` state. When the End Of Media marker is reached during a tape write operation, the volume will enter `EOM` state. If the number of segments on the volume drops to zero after reaching EOM, the state will change to `Empty`. The volume may be placed in `Deny` state by the server depending on the setting of the VV Condition.

Storage Class

The storage class to which the tape virtual volume is assigned.

Map Flags

Information about the state of the volume. Contains one of the following values:

- `Never Written` indicates data has not been written to the volume.
- `Retired` indicates new data will not be written to this volume.
- `Repack Output` indicates all files on this volume were written there by the **repack** process.

Max Written Length

The total number of bytes that have been written to the volume. Each time the volume is written, this number is incremented by the number of bytes written. This number is not decremented when storage segments are deleted from the volume. This value does not include parity bytes in RAIT volumes.

Active Length

The sum of the number of bytes in active storage segments. As segments are added to the volume, this number increases. As segments are deleted, this number decreases. This value does not include parity bytes in RAIT volumes.

Estimated Size

The estimated size of the volume in bytes. It is only an estimate of the number of bytes that can be written on a tape volume. Depending on the compression characteristics of the data, the actual number of bytes may be smaller or greater. This value does not include parity bytes in RAIT volumes.

Space Left

The estimated number of bytes that may still be written to the tape virtual volume. It is initially set to **Estimated Size** and is decremented as storage segments are written. It is not incremented when storage segments are deleted. This value does not include parity bytes in RAIT volumes.

Because of the variable compressibility of data written to tapes, this value may behave strangely. For instance, the end of the media (EOM) may be reached earlier than estimated, causing this value to become zero quite suddenly. Conversely, EOM may be reached later than expected. In this case, **Space Left** will be decremented until it reaches a value of "1", then remain at "1" until EOM is reached.

Number of Reads

The number of times the tape volume has been read.

Number of Writes

The number of times the tape volume has been written.

Time Last Read

The date and time the volume was last read. If this field is blank, the tape volume has not been read.

Time Last Written

The date and time the volume was last written. If this field is blank, the tape volume has not been written.

Time Last Verified

The date and time the volume was last verified (via the **lbp_verify** tool). If this field is blank, the tape volume has never been verified.

PVL Job ID

The PVL job ID in which this volume is mounted. This field will be zero if the volume is not mounted.

Current Position

The address to which the tape is positioned, expressed as an HPSS Relative Stripe Address. The Relative Address consists of two parts, the tape section and the byte offset. Tape sections are the data written between a pair of tape marks. The byte offset is the offset from the beginning of the current section to the place where the tape read or write is to take place. When reading the tape, this address is set to the initial location where the tape read begins, then advances as the read progresses. When writing a tape, this value is set to the **Next Write Address**, and then advances as the tape is written. Updates to this field occur when I/O operations complete, so in some cases significant amounts of time may pass between updates.

Next Write Address

The next address that will be written on the tape volume, expressed as an HPSS Relative Stripe Address.

Core Server Tape Volume

File Edit Help

Name: JC067200

VV Condition: RWC ☐ Retired

Changes Pending: None ☐ FIFO

More commonly used data Less commonly used data RAIT Options

File Family: Not in a family

VV BlockSize: 1048576

PV BlockSize: 262144

Max Blocks Between Tape Marks: 43200

Max Blocks Between Unbuffered Tape Marks: 43200

Data Stripe Width: 1

Aggregates: 5

File Segments in Tape Aggregates: 48713

Bytes in Aggregates: 91,559,056

Estimated Files on VV: 48717

Creation Time: Jan 24, 2022 3:53:36.000 PM

☐ Block Protection

VV ID:

Object Id: 00000539-04-00000003-01ec7d6015353ea6-0000

Current Writing Segment:

Object Id: 00000000-00-00000000-0000000000000000-0000

Physical Volumes

Vol Name ▼	Type	Dev ID	Condition	Errors	Format	Mvr IP Addr	Mvr	Mvr Host
JC067200	IBM - 3592 E07/EH7 (TS1140) JC/JY Tape	0	RW		HPSS 0.0.0.0[0]			

1/1

☐ Freeze Update Refresh Dismiss

Retrieving data - Succeeded

Less commonly used data tab

File Family

The family to which the volume is assigned, if any. If it is not assigned to a family, it is assigned to the default family, family zero.

VV Block Size

The virtual volume block size. This is the number of bytes written from a data stream to an element of the striped volume before the stream switches to the next element of the stripe.

PV Block Size

The size, in bytes, of the media data block.

Max Blocks Between Tape Marks

The maximum number of media blocks that will be written to this tape before a tape mark is written. The product of this value and the Media Block Size defines the Physical Volume Section Length.

Data Stripe Width

The number of physical volumes in the stripe group that contain file data. All volumes in HPSS are considered to be striped, even if this value is one.

Aggregates

The number of tape aggregates on this volume.

File Segments in Tape Aggregates

The number of file segments in the tape aggregates on this volume. This is a count of files within all aggregates on the volume. Files may be split across tape aggregates on the volume, which can lead them to be counted multiple times in this statistic.

Bytes in Aggregates

The number of bytes stored in the tape aggregates on this volume. This value does not include parity bytes in RAIT volumes.

Estimated Files on VV

An estimate of the number of files stored on the tape volume. This estimate is arrived at by summing the number of individual tape storage segments on the volume (the server assumes one file per segment) and adding the number of files stored in aggregates on the volume. Since the number of files in an aggregate can be relatively large, the number of files on the volume may be much larger than **Number of Active Segments**.

Creation Time

The date and time the tape volume was created in the Core Server.

Block Protection

When this flag is set, Logical Block Protect (LBP) will be enabled for the physical tape volumes in this virtual tape volume. When LBP is enabled, HPSS will generate and pass a CRC with each block written to the tape volume and validate the CRC accompanying each block read from the tape volume.

VV ID

The detailed description of the tape virtual volume ID associated with this tape virtual volume.

Current Writing Segment

The detailed description of the tape storage segment ID associated with the segment that is found at the end of the tape volume. For various reasons, some tape volumes do not have a value in this field. It will be displayed filled with blanks in that case.

The screenshot shows the 'Core Server Tape Volume' window with the 'RAIT Options' tab selected. The window has a menu bar with 'File', 'Edit', and 'Help'. The main configuration area includes fields for 'Name' (X0003000), 'VV Condition' (RWC), 'Retired' (unchecked), 'Changes Pending' (None), and 'FIFO' (checked). Below these are three tabs: 'More commonly used data', 'Less commonly used data', and 'RAIT Options'. The 'RAIT Options' tab contains fields for 'Parity Stripe Width' (0), 'Minimum Write Parity' (0), 'Number of RAIT Engines' (0), 'Read Verification' (unchecked), and 'Mount Minimum PVs' (unchecked). At the bottom, there is a 'Physical Volumes' table with columns: Vol Name, Type, Dev ID, Condition, Errors, and Format. The table shows one entry: X0003000, Generic - Default Tape, 0, RW, HPS. Below the table are buttons for 'Freeze', 'Update', 'Refresh', and 'Dismiss', and a status bar showing 'Retrieving data - Succeeded'.

Name: X0003000

VV Condition: RWC ☐ Retired

Changes Pending: None ☒ FIFO

More commonly used data Less commonly used data **RAIT Options**

Parity Stripe Width: 0 (0 = RAIT is disabled)

Minimum Write Parity: 0

Number of RAIT Engines: 0

☐ Read Verification ☐ Mount Minimum PVs

Physical Volumes

Vol Name ▼	Type	Dev ID	Condition	Errors	Format
X0003000	Generic - Default Tape	0	RW		HPS

1/1

☐ Freeze Update Refresh Dismiss

Retrieving data - Succeeded

RAIT Options tab

Parity Stripe Width

The number of physical volumes in the tape virtual volume that contain parity information. If this value is zero, the VV is not RAIT.

Minimum Write Parity

The minimum number of parity PVs that must remain writable in order for the RAIT VV to be writable.

Number of RAIT Engines

The number of RAIT Engines that will be used when reading or writing the volume.

Read Verification

When set, parity information in the RAIT VV will be used to verify the integrity of data read from the VV. This field can be modified on this screen.

Mount Minimum PVs

When set, the smallest number of PVs necessary will be mounted when reading the RAIT VV. When **Read Verification** is not set, this will be equal to the value of DSW; and when **Read Verification** is set, this will equal (DSW + 1). This field can be modified on this screen.

Physical Volumes

This is a table of physical volume attributes for the physical volumes that make up this tape virtual volume.

Vol Name

The ASCII name of the physical volume.

Type

The media type.

Dev ID

The ID of the device the physical volume is mounted on. If the volume is not mounted, this field will be zero.

Condition

The administrative condition of the individual PV. In non-RAIT VVs, this value will always be RW. In RAIT VVs, this value may be RW, RO, or DOWN and may be changed using this screen.

Errors

The accumulated error history of the PV. In non-RAIT VVs, this field will always be blank. In RAIT VVs, this field may contain an M if the PV had a mount error, an R if the PV had a read error, and a W if the PV had a write error. The specific error history is recorded in metadata and accessible via the **lspvhist** utility program.

Format

The recording format of the tape. Native HPSS tapes will show HPSS in this field. Tapes imported to HPSS from other systems will show other values.

Mvr IP Addr

The IP address of the Mover that operates this physical volume. If the volume is not mounted, this field will be zeros.

Mvr

The descriptive name of the Mover that operates this physical volume. If the volume is not mounted, this field will be blank.

Mvr Host

The name of the host on which the Mover runs. If the volume is not mounted, this field will be blank.

8.5.5. Changing Core Server volume condition

HPSS provides a single control that sets the operational condition of a disk or tape virtual volume. The control is found on the *Core Server Tape Volume* and *Core Server Disk Volume* windows and is called the VV Condition.

VV Condition controls the availability of the volume for the following actions:

- Creation of new storage segments
- Reading of existing storage segments
- Writing of existing storage segments
- Mounting of tape media

Tape volumes have five possible settings for VV Condition:

RWC

Read, Write, Create

RO

Read Only

EOM

End of Media reached

EMPTY

The volume reached EOM, and all data has been removed.

DOWN

Volume not available

Disk volumes have five possible settings for VV Condition:

RWC

Read, Write, Create

RW

Read, Write

RO

Read Only

EMPTY

The volume is in RO condition and all data has been removed.

DOWN

Volume not available

A volume in RWC condition is fully enabled for use. This is the normal setting for all disk volumes and for tape volumes that have not reached EOM. Storage segments can be created on the volume;

existing segments can be read and written. Volumes in RWC condition can be changed to any other condition.

Existing storage segments on disk volumes in RW condition can be read and written, but new segments cannot be created. This setting is useful when emptying a disk through attrition. Disk volumes can be emptied more quickly by setting the condition to RW, then purging or repacking the disk. Volumes in RW condition can be changed to any other condition except as limited by the Retire flag (see below).

RO condition is similar to RW condition, but writes to existing disk storage segments are not permitted.

Storage segments on tape volumes in RO condition can be read, but not written. New segments cannot be created. Tapes in RO condition are similar to tapes in EOM condition with the exception that their condition can be changed to RWC or RW while tapes in EOM condition cannot. Volumes in RO condition can be changed to any other condition except as limited by the Retire flag (see below).

Tape volumes enter EOM condition when the end of media marker is reached while writing the tape or when the administrator sets the condition to EOM. EOM condition is similar to RO condition - segments can be read, but not written. Unlike RO condition, tapes in EOM condition can only be changed to DOWN. EOM volumes cannot enter either RWC or RO condition.

Volumes in DOWN condition cannot be read, written, created on or mounted. This setting effectively removes the volume from the system while maintaining the records of the storage segments on it and is useful for dealing with failed disk or tape volumes. Disks in DOWN condition can be changed to RWC, RW or RO condition. Tapes in DOWN condition can be changed to RWC or RO condition if they have not reached EOM. Tapes that have reached EOM can only be changed from DOWN to EOM.

Disk or tape volumes may be reported in EMPTY condition, but the administrator can never change a disk or tape condition to EMPTY manually. This condition will be reported when a tape in EOM condition becomes empty, or a disk in RO condition becomes empty.

The **Retire** flag controls entry into RWC condition. When the **Retire** flag is set, the condition cannot be changed to RWC, but all other changes are permitted subject to the rules outlined above. By preventing entry to RWC condition, retired volumes can be read and written, but new storage segments cannot be added to the volume. Normal attrition will gradually empty the volume.

To retire a disk or tape volume, set the condition to any setting other than RWC (typically RW for disks and RO for tapes) and then set the **Retire** flag. These two changes can be done in a single update of the *Core Server Disk Volume* or *Core Server Tape Volume* window. Note that tapes do not need to be in EOM condition to be retired.

The **Retire** flag can be cleared if necessary. Once cleared, the condition can be changed to RWC, subject to the rules outlined above.

8.5.6. Moving PVR cartridges to another PVR

The intent of the Move Cartridge operation is to notify HPSS that a cartridge has been moved from one tape library to another.

Note that this operation will not cause a cartridge to move. Instead, the cartridge should be moved from the old library to the new library manually before informing HPSS. The operation will fail if the cartridge has not been ejected from its original robot and injected into the new robot. See the operator manuals for the specific robots involved to determine the procedures for manually ejecting and injecting a cartridge. Either the source or destination PVR (or both) during a move operation may be an operator mounted set of drives instead of a robot.

Specifically for shelf tapes, when a Move Cartridge is issued to a shelf tape, the target PVR will check whether or not the cartridge is physically in the target library. If it is not, then it will be treated as a move of a shelf tape to a shelf tape in the target library (that is, maintain its checked-out/shelf status). If it is physically in the target library, then it will be treated as a move of a shelf tape to an online cartridge in the target library.

Use the *Move Cartridges to New PVR* window in the following section.

8.5.6.1. Move Cartridges window

This window allows you to change which PVR owns a set of cartridges. Before initiating the request from the SSM window, the cartridges must already be physically placed into a tape library managed by the new PVR.

The list of the volumes to be moved may be constructed in any of three ways. Each volume name may be typed in one at a time, or a list of volume names may be automatically generated from a single entry, or a list of volume names may be specified from an input file. Volume names are not entered directly into the **Volume List** at the bottom of the window.

Any of the three entry methods may be repeated multiple times on the same window to add additional volumes to the list. All three entry methods or any combination of them may be used in succession on the same window.

To add a single volume name to the **Volume List**, set the **Fill Count** and the **Fill Increment** fields each to "1" and type the volume name into the **Volume Label** field. The volume name will be added to the **Volume List**.

To automatically generate a list of volume names in the **Volume List**, set the **Fill Count** to the desired number of volumes. Set the **Fill Increment** to the number by which each automatically generated label should differ from the previous one. Then type the starting volume name into the **Volume Label** field. The specified number of volume names will be added to the **Volume List**, each one larger than the previous entry by the specified **Fill Increment**.

Example:

```
Fill Count = 6
Fill Increment = 10
Volume Label = "AA0070"

Labels automatically inserted into Volume List:
"AA0070"
"AA0080"
"AA0090"
"AA0100"
"AA0110"
"AA0120"
```

When an addition produces overflow in a column, numerical columns are carried over properly to alphabetic columns and vice versa.

Example:

```
Fill Count = 6
Fill Increment = 2000
Volume Label= "AA7329"

Labels automatically inserted into Volume List:
"AA7329"
"AA9329"
"AB1329"
"AB3329"
"AB5329"
"AB7329"
```

The filling will not occur and an error will be displayed if the specified values would generate an invalid volume label (for example, one greater than zzz999).

To specify a list of volumes from a file, create a file containing the name of each volume to be moved on a separate line. Volume names must be six alphanumeric characters. No other characters are allowed on the line. The file must be accessible from the host on which the hpssgui is executing. Enter the name of the file in the **File Containing Volume List** field. The volume names from the file will be added to the **Volume List**.

Each cartridge in the list is sent to the System Manager as a separate move request. If an error occurs on a request, SSM stops processing the list of cartridges at that point. The window status line will

report the number of successful moves, and the cartridge name that caused the failure. For example, if you try to move a cartridge to a PVR, and that cartridge is already controlled by that PVR, an error will occur. Therefore, if a multi-cartridge move fails, you cannot click the **Move Cartridge** button to retry the same list of cartridges, because the ones that succeeded before will now fail. You must first remove the successfully moved cartridges from the list. The **Clear Selected** button is provided for this purpose.

Field descriptions.

New PVR

The descriptive name of the PVR to which the cartridges will be moved.

File Containing Volume List

The name of an external file containing a list of cartridge labels to be added to the end of the **Volume List**.

Fill Count

The number of cartridge labels to be added to the end of the list when the **Volume Label** field is next modified. This number may be one or greater. If the list fills up before the **Fill Count** is exhausted, filling stops, and a message box is displayed (see **Maximum Volumes Allowed** below).

Fill Increment

In a multiple-cartridge fill, this field determines how each new cartridge label is generated from the previous one. A cartridge label is six alphanumeric characters. The **Fill Increment** is added to the least significant part of a cartridge label to generate a new label.

Volume Label

The cartridge label to be added to the end of the **Volume List**. If **Fill Count** is greater than "1", this will be the first label of a fill sequence. The label must be exactly six characters long.

Maximum Volumes Allowed

The maximum number of cartridge labels that will be allowed in the **Volume List** (10,000).

Total Count

The total number of labels currently in the **Volume List**.

Volume List

A list of cartridge/volume labels specifying the cartridges to be moved. You cannot enter labels directly into this list, but must construct it in one of the three ways described above. Use the scrollbar to move up and down the list of labels.

You can select one or more labels from the list for the purpose of removing them from the list (see **Clear Selected** below).

Buttons.

Clear All

Clears the **Volume List**, and resets **Fill Count** and **Fill Increment** to "1".

Clear Selected

If one or more cartridges are selected (highlighted) in the **Volume List**, clicking on this button will remove them from the list. Note that this does not actually delete anything from HPSS.

Move Cartridges

Begins moving the listed cartridges. A start message is displayed on the message line at the bottom of the window, and all window features are disabled, except for the **Dismiss** button; this prevents you from modifying any data on the window while the move is in progress.

It is a good idea to have the *HPSS Alarms and Events* window open while moving proceeds, so that all relevant log messages can be viewed. Move operation initiation, completion and errors will be reported in the *HPSS Alarms and Events* window.

While moving proceeds, status messages are displayed on the status bar. When the move is finished, a completion message is displayed and the window features are re-enabled. At this point new data can be entered for a new move.

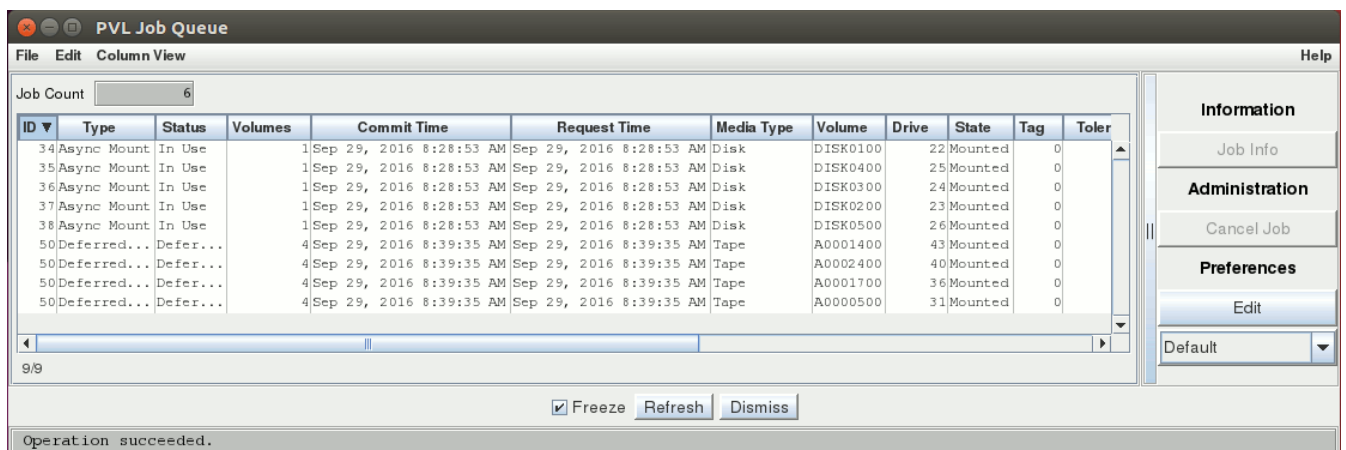
8.6. Monitoring and managing volume mounts

Volume mounts are performed by the PVL. Every volume mount is associated with a PVL job. The *PVL Job Queue* window lists all PVL jobs and volume mounts for each job. The *PVL Request Information* window lists detailed information about a single PVL job.

The PVL reports every tape mount request in the *Tape Mount Requests* window. Tapes which need to be returned to the robot from a shelf or vault location are reported in the *Tape Check-In Requests* window.

Each of these windows is described in detail in this section.

8.6.1. PVL Job Queue window



This window shows all outstanding jobs in the PVL and the mounted volume information for each job. From this window, the user can issue a request to view more information for a particular PVL job or to cancel it. Each PVL job represents a volume mount (or series of mounts for a striped disk or tape volume).

Field descriptions.

Job Count

The number of PVL jobs in the queue.

Job List

This is the main part of the window, consisting of a table of job information, a title line containing labels for each column, and vertical and horizontal scrollbars.

The function buttons to the right of the list require that a job be selected from the list. To select a job, click on it with the mouse; the selection will be highlighted.

The fields displayed in the table, as identified in the title line column headings, are shown below.

ID

A unique number assigned to each job. There may be multiple entries per job ID.

Type

The type of job. Possible types are:

- `Async Mount` indicates an asynchronous mount.
- `Default Import` indicates a media import of type "Default Import".
- `Scratch Import` indicates a media import of type "Scratch Import".
- `Overwrite Import` indicates a media import of type "Overwrite Import".
- `Export` indicates a media export.
- `Deferred Dismount` indicates a cartridge is scheduled for dismount, but the dismount has been delayed in case the cartridge needs to be used by another upcoming job.
- `Move` indicates a cartridge being moved to a new PVR.
- `Relabel` indicates a cartridge being relabeled.
- `Sync Mount` indicates a synchronous mount.
- `Tape Check-In` indicates a cartridge being added to the library.
- `Tape Check-Out` indicates a cartridge being removed from the library to be placed on the shelf or vault.
- `In Use Reduced` indicates a mount error occurred but the resulting set of mounted PV's are within the minimum number allowed for a RAIT volume.

Status

The current status of the job. Possible values are:

- `Uncommitted` indicates the mount job is a multi-part job that has been started, but the last volumes have not been added, and the mount operation has not been committed.
- `Cartridge Wait` indicates the job is waiting for another job to release a cartridge that it needs.
- `Drive Wait` indicates the job is waiting for a drive to become available.

- `Mount Wait` indicates the job is waiting for a volume to be mounted.
- `Mounted` indicates all volumes required for the job are mounted.
- `Dismount Pending` indicates volumes are in the process of being dismounted.
- `Aborting` indicates the job is being aborted.
- `Inject` indicates cartridges are being injected into a PVR (for import jobs only).
- `Eject` indicates cartridges are being ejected from a PVR (for export jobs only).
- `In Use` indicates all volumes required for the job are mounted and ready for use.
- `Deferred Dismount` indicates dismount for cartridges will be delayed.
- `Tape Check-In` indicates the job is waiting for the cartridge to be checked in.
- `Tape Check-Out` indicates the job is waiting for the cartridge to be checked out.
- `Completed` indicates the job is complete.

Volumes

The number of volumes that are mounted for this job.

Commit Time

The date and time when the job was committed. This will be the same as **Request Time** for all jobs except asynchronous mounts.

Request Time

The date and time when the job was sent to the PVL

Media Type

Type of media being used by job. `Disk`, `Tape`, or `Unknown`. This field is useful for filtering the **Job List**.

Volume

The eight-character volume label.

Drive

The numeric identifier of the drive in which the volume is mounted, or `none` if the volume does not have a drive associated with it.

State

The state of the volume within the context of this request. Possible values are:

- `Cart Assigned`
- `Cart Wait`
- `Dismount Pending`
- `Dismounted`

- Drive Wait
- Eject
- Inject
- Mount Failed
- Mount Pending
- Mounted
- Reading Label
- Tape Check-In
- Tape Check-Out
- Uncommitted
- Unload Pending

Tag

A value used to associate the physical volumes within a RAIT volume; that is, all tape volumes in a RAIT volume will have identical **Tag** values.

Tolerate

Specifies the number of physical tape volumes that can fail mounting within a RAIT volume (identical **Tag** values) and still be considered successfully mounted.

Library Unit ID (tape only)

Since library architecture varies, the method to associate related drives at a unit level will be determined through the use of a "unit ID" that is placed in the drive metadata. This will allow the administrator to analyze their specific library configuration and apply their strategy as they feel fit. The use of "unit ID" allows the library specifics to be abstracted from the PVL. Library details are the domain of the PVR. See **Defer Dismount Exempt Count** information in *Section 5.2.7.3, "SCSI PVR-specific configuration window"*.

Drive Type

The HPSS drive type assigned to the drive.

Mover

The name of the Mover that owns the device where the volume is mounted. This field will be blank if the volume is not currently mounted.

PVR

The name of the PVR that owns the mounted volume. This field will be blank if the volume is not a tape volume.

Buttons.**Job Info**

Click on this button to open the *PVL Request Information* window for the selected job.

Cancel Job

Click on this button to cancel the selected job. You are prompted to confirm your request before it is actually executed. This will generate a dismount request.

Related Information.

Section 8.6.5, "Tape Mount Requests window"

Section 8.6.3, "Canceling queued PVL requests"

8.6.2. PVL Request Information window

The screenshot shows the 'PVL Request Information' window. It has a menu bar with 'File', 'Edit', and 'Help'. The main area contains the following fields:

- Job ID: 50
- Request Type: Deferred Dismount
- Request Timestamp: Sep 29, 2016 8:39:35 AM
- Mounted Volumes: 4
- Request Status: Deferred Dismount

Below these fields is a table with the following columns: Volume, Drive, State, Tag, Tolerate, Library Unit ID, Drive Type, Mover, and PVR.

Volume	Drive	State	Tag	Tolerate	Library Unit ID	Drive Type	Mover	PVR
A0000500	31	Mounted	0	0	-1	STK - T10000B Tape	Mover (hpss-dev-dus03)	SCSI PVR
A0001400	43	Mounted	0	0	-1	STK - T10000B Tape	Mover (hpss-dev-dus03)	SCSI PVR
A0001700	36	Mounted	0	0	-1	STK - T10000B Tape	Mover (hpss-dev-dus03)	SCSI PVR
A0002400	40	Mounted	0	0	-1	STK - T10000B Tape	Mover (hpss-dev-dus03)	SCSI PVR

At the bottom of the table, there is a status bar showing '4/4'. Below the table are three buttons: 'Freeze' (checked), 'Refresh', and 'Dismiss'. At the very bottom, a status bar says 'Retrieving data - Succeeded'.

This window is displayed when the **Job Info** button is clicked on the *PVL Job Queue* window. It allows you to view the information associated with a PVL job/request.

Field descriptions.

Job ID

The unique number assigned to the job being viewed.

Request Type

The type of job/request. Possible types are:

- Async Mount indicates an asynchronous mount.
- Default Import indicates a media import of type "Default Import".
- Scratch Import indicates a media import of type "Scratch Import".
- Overwrite Import indicates a media import of type "Overwrite Import".
- Export indicates a media export.
- Move indicates a cartridge being moved to a new PVR.

- `Relabel` indicates a cartridge being relabeled.
- `Sync Mount` indicates a synchronous mount.
- `Deferred Dismount` indicates dismount delayed.
- `Tape Check-In` indicates a cartridge being added to the library.
- `Tape Check-Out` indicates a cartridge being removed from the library.
- `In Use Reduced` indicates a mount error occurred but the resulting set of mounted PVs are within the minimum number allowed for a RAIT volume.

Request Timestamp

The date and time when the request was issued to the PVL.

Mounted Volumes

The number of volumes that are currently mounted for this request.

Request Status

The current status of the job/request. Possible values are:

- `Aborting` indicates the job is being aborted.
- `Cartridge Wait` indicates the job is waiting for another job to release a cartridge that it needs.
- `Completed` indicates the job is completed. Once a job is completed, it no longer exists in the PVL job queue, and this window will no longer receive any updates.
- `Deferred Dismount` indicates dismount for cartridges will be delayed.
- `Dismount Pending` indicates volumes are in the process of being dismounted.
- `Drive Wait` indicates the job is waiting for a drive to become available.
- `Eject` indicates cartridges are being ejected from a PVR (for export jobs only).
- `In Use` indicates all volumes required for the job are mounted and ready for use.
- `Inject` indicates cartridges are being injected into a PVR (for import jobs only).
- `Mount Wait` indicates the job is waiting for a volume to be mounted.
- `Mounted` indicates all volumes required for the job are mounted.
- `Tape Check-In` indicates the job is waiting for the cartridge to be checked in.
- `Tape Check-Out` indicates the job is waiting for the cartridge to be checked out.
- `Uncommitted` indicates the job has not yet been committed (used for asynchronous mount jobs only).

Volume List

This group of fields shows information on the volumes associated with this request.

Volume

The eight-character volume label.

Drive

The numeric identifier of the drive in which the volume is mounted, or `none` if the volume does not have a drive associated with it.

State

The state of the volume within the context of this request. Possible values are:

- `Cart Assigned`
- `Cart Wait`
- `Dismount Pending`
- `Dismounted`
- `Drive Wait`
- `Eject`
- `Inject`
- `Mount Failed`
- `Mount Pending`
- `Mounted`
- `Reading Label`
- `Tape Check-In`
- `Tape Check-Out`
- `Uncommitted`
- `Unload Pending`

Tag

A value used to associate the physical volumes within a RAIT volume; that is, all tape volumes in a RAIT volume will have identical **Tag** values.

Tolerate

Specifies the number of physical tape volumes that can fail mounting within a RAIT volume (identical **Tag** values) and still be considered successfully mounted.

Library Unit ID (tape only)

Since library architecture varies, the method to associate related drives at a unit level will be determined through the use of a "unit ID" that is placed in the drive metadata. This will allow the administrator to analyze their specific library configuration and apply their strategy as they feel fit. The use of "unit ID" allows the library specifics to be abstracted from the PVL.

Library details are the domain of the PVR. See **Defer Dismount Exempt Count** information in *Section 5.2.7.3, “SCSI PVR-specific configuration window”*.

Drive Type

The HPSS drive type assigned to the drive.

Mover

The name of the Mover that owns the device where the volume is mounted. This field will be blank if the volume is not currently mounted.

PVR

The name of the PVR that owns the mounted volume. This field will be blank if the volume is not a tape volume.

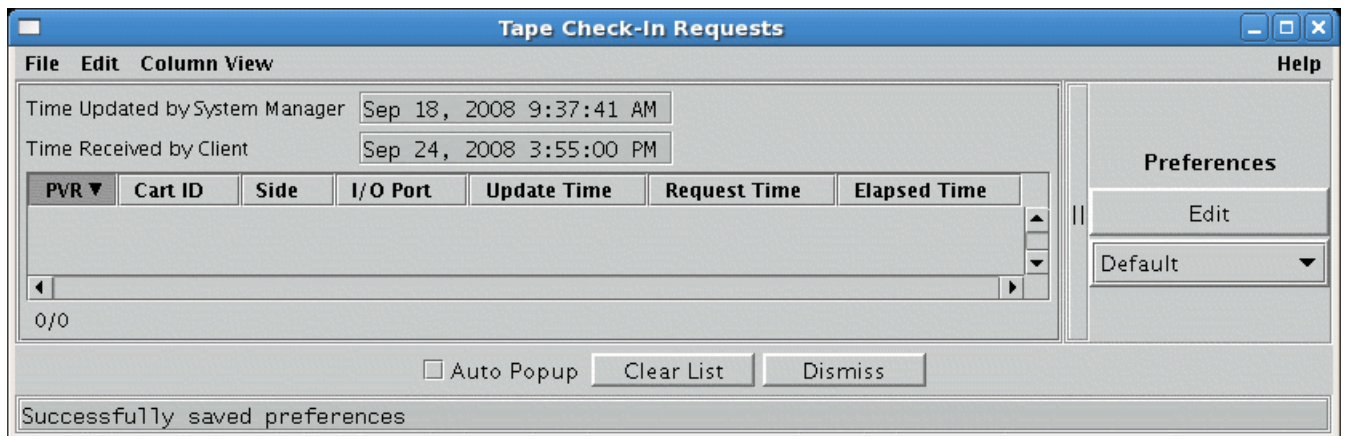
8.6.3. Canceling queued PVL requests

PVL requests that cannot be completed due to a hardware or software problem that HPSS cannot handle can be canceled by the system administrator.

Canceling a PVL job will result in the PVL issuing a dismount job request if any of the volumes in the job were mounted. The dismount request causes an unload request to be sent to the controlling Mover. This request may fail if the device is performing an I/O transfer. The PVL will retry the unload command until the I/O completes and the device reservation is freed. Once a volume has been successfully unloaded, all I/O operations directed to the device will fail. Cancellation should therefore be done only after careful consideration of its consequences.

When a deferred dismount job is canceled, the PVL issues dismount requests to all drives loaded with cartridges in deferred dismount state. Canceling a deferred dismount job can be done as often as necessary and does not result in any errors elsewhere in the system.

8.6.4. Tape Check-In Requests window



When using the shelf-tape feature of HPSS, operators are notified via the *Tape Check-In Requests* window of mount requests for tapes which are stored outside the robot (such as in a vault). Once the requested tapes are inserted back into the robot, tape check-in requests are removed from the list. If a Tape Check-In Request has been displayed, and the corresponding tape has not been checked in, the PVR will begin logging alarms indicating that the appropriate tape has yet to be checked in. The

frequency of these alarms is controlled by the Shelf Tape Check-In Alarm field of the PVR-specific configuration window.

The requests are displayed in chronological order. Each time such a request is received by SSM, it is added to the list, but duplicate requests for the same tape are not displayed. When a message is received indicating that the check-in request has been satisfied, the request is removed from the window. For operator PVRs such check-in requests mean that a tape must be mounted by hand.

The maximum number of tape check-in messages that can be displayed on the *Tape Check-In Requests* window is 100. When this limit is exceeded, some tape check-in notifications will be lost.

If the **Auto Popup** check box is selected, this window will automatically reopen whenever a new tape check-in request is received. To open this window manually, select the **Monitor** menu on the *HPSS Health and Status* window, from there select the **Tape Requests** submenu, and from there select the **Check-In** menu item.

The HPSS **shelf_tape** utility may be used to move tapes from offline storage to a tape library.

Field descriptions.

Tape Check-In Requests list

PVR

The descriptive name of the PVR that owns the cartridge.

Cart ID

The six-character label identifying the cartridge to be inserted into the robot.

Side

The side (partition) number to be mounted. This may be blank.

I/O Port

The descriptive name of the I/O port where the cartridge will be placed.

Update Time

The date and time when the check-in request was last updated.

Request Time

The date and time that the check-in request was first received by the SSM.

Elapsed Time

The length of time (days-hours:min:sec) that the check-in request has been waiting.

Buttons and check boxes.

Auto Popup

If ON, this window will be opened automatically if it is not already on the screen when a check-in request is received.

Clear List

Clears the list of check-in requests. Note that this does not cancel any check-in requests, but just removes them from the list. Pending check-in requests will reappear in the window as the PVR periodically retries the check-in of cartridges. This can be useful for removing stale check-in

requests that, for some reason, never issued a completion message to SSM. When this button is clicked from any SSM session, the *Tape Check-In Requests* windows on all SSM sessions will be cleared.

8.6.5. Tape Mount Requests window



The *Tape Mount Requests* window displays tapes which need to be mounted in a drive. All HPSS tape mount requests, including both robotic and operator tape mounts, will be displayed in the window. For operator PVRs, such mount requests mean that a tape must be mounted by hand. When mount requests for robotic PVRs do not disappear from the window in a timely manner, it can be an indication of hardware or other problem in the robot.

Outstanding tape mounts, along with the associated PVR, are displayed in chronological order. Each time a tape mount request is received by SSM, it is added to the list, but duplicate requests for the same cartridge are not displayed. When a message is received indicating that the mount request has been satisfied, the request is removed from the window.

The window may display an old mount request at times, due to lost notifications. Click on the **Clear List** button to refresh the window.

When a new mount notification is received, SSM may bring up the *Tape Mount Requests* window automatically if it is not currently displayed. This behavior is disabled by deselecting the **Auto Popup** check box. To open the window manually, select the **Monitor** menu on the *HPSS Health and Status* window, from there select the **Tape Requests** submenu, and from there select the **Mount** menu item.

Field descriptions.

Tape Mount Requests list

PVR

The descriptive name of the PVR that owns the cartridge.

Cart ID

The six-character label identifying the cartridge to be mounted.

Side

The side (partition) number to be mounted. This may be blank.

Drive

The drive address if a specific drive was requested for the mount. In most cases, a drive is not specified, and this is blank.

Update Time

The date and time when the mount request was last updated.

Request Time

The date and time that the mount request was first received by the SSM.

Elapsed Time

The length of time (days-hours:min:sec) that the mount request has been waiting.

Buttons and check boxes.**Auto Popup**

If ON, this window will be opened automatically if it is not already on the screen when a mount request is received.

Clear List

Clears the list of mount requests. Note that this does not cancel any mount requests, but just removes them from the list. Pending mounts will reappear in the window as the PVR periodically retries the mounts. This can be useful for removing stale mount requests that, for some reason, never issued a completion message to SSM. When this button is clicked from any SSM session, the *Tape Mount Requests* windows on all SSM sessions will be cleared.

8.6.6. Administrative tape dismounts

HPSS provides the ability to administratively command a tape dismount in unusual situations, such as when an administrator wants a tape dismounted without waiting for HPSS to do so, or when HPSS is unable to do so due to a failure.

Administrative dismounts should only occur in exceptional error circumstances. Dismounting a tape may cause the operations being performed on that tape to fail. A library dismount operation is unlikely to occur while a Mover has the device reserved for I/O.

Use the *Devices and Drives* list window to issue a dismount for the desired device/drive entry. A confirmation window will pop up requiring you to confirm the dismount request.

8.7. Storage technology replacement

As next generation storage technology becomes available, the need arises to replace prior generation storage media. The procedure for effecting this technology change involves updating a storage class's definition with the new media type and characteristics, importing new media, creating new virtual volumes (VVs), retiring prior generation VVs, moving data from old media to new media, and finally removing the old media altogether from HPSS.

The act of retiring a VV indicates to HPSS that no more data should be written to the volume. Furthermore, it indicates that the volume should not be reclaimed for future reuse. A retired volume is considered end of life, only to be used for light reading until all the data is moved off the volume. Note that a retired VV cannot have a **VV Condition** of RWC.

To replace prior generation media with new, the following procedure serves as a guide. HPSS administrators should be aware that the procedure will probably need to be tailored to each customer's specific circumstances and needs. This procedure applies to both tape and disk volumes.

1. Change the storage class definition by changing its media type and any other related fields.

- The new storage class values must reflect the characteristics of the new media.
- In addition, the new values should be compatible with other storage levels in the hierarchy.



See *Section 6.1.5, "Changing a storage class definition"* for guidance on changing a storage class's **Media Type** and other values.

Changing a storage class's configuration does not apply to existent HPSS virtual volumes; a VV's storage class configuration is set upon resource creation and remains immutable for the lifetime of the VV. A new storage class configuration will only apply to new VVs that are created after modifying the storage class.



If upgrading to a new subtype of the same tape type, setting HPSS environment variable `HPSS_CORE_CREATE_ALLOW_SUBTYPE` to "true" will allow you to import the new tape subtypes without first changing the storage class configuration. This is especially useful when creating resources of multiple different subtypes at once. Regardless, the storage class configuration should still be changed to the correct subtype for future tape VV creation. This environment variable is only provided to ease the transition.

2. Click **Apply Config** on the *HPSS Health and Status* screen for the modifications to the storage class to take effect.

3. Import and create new volumes.

- See *Section 8.1, "Adding storage space"* for details.
- The new volumes will take on the characteristics of the redefined storage class.
- Consider creating only one volume at first in order to check its characteristics. If they are not as desired, delete the volume and repeat the first two steps in this procedure.



It isn't necessary to create all the new volumes at once. Only as many as are required to meet the immediate demands of the system need to be created. Keep in mind that, when storage resources are created, they are immediately available for use by HPSS.

4. Use the **retire** utility program to retire the prior old volumes.

- View the **retire** man page for details on how to use the program.



dump_sspvs can be used to create a list of physical volumes for **retire** to process. Alternatively, if the number of volumes to retire is small, the **hpssgui** may be used to retire them one-by-one. Do so by changing each volumes' **VV Condition** and selecting the *Retire* check box in the *Core Server Tape Volume* or *Core Server Disk Volume* window. It is not necessary to change the **VV Condition** of tape volumes to EOM; the RO VV Condition will prevent the tapes from being written and is reversible (unlike EOM).

5. Use the **repack** utility program to move data from all retired volumes in the storage class.
 - This step is frequently a long process and may require regular attention from the HPSS administrators.
 - Use **dump_sspvs** to monitor the **repack** progress.
 - View the **repack** man page for details on how to use the program.
6. Use the **remove** utility program to remove all retired and empty volumes.
 - This program will delete all resources associated with these volumes and export them from HPSS.
 - View the **remove** man page for details on how to use the program.

Chapter 9. Logging and status

9.1. Logging overview

The purpose of logging is to record events of interest that occur in HPSS in the sequence they occur to support diagnostic research.

HPSS provides eight log message types:

- Alarm
- Event
- Info
- Debug
- Request
- Security
- Accounting
- Trace

The purpose of each of these log message types is described later in this chapter.

The HPSS logging system includes these components:

- Log policies (*Section 9.2, “Log policies”*)
- Syslog (*Section 9.3, “Managing syslog”*)
- Log archiving (*Section 9.3.9, “Managing log archiving”*)
- The SSM *Alarms and Events* window (*Section 9.4.1, “Alarms and Events window”*)

A standard configuration for logging services is usually set by the administrator during the HPSS system configuration. Specialized configurations can be set up and used to provide more or less logging for site-specific or shift-specific operational situations. Increasing the level of logging may slow down overall system operation due to the overhead of extra messages, and decreasing the amount of logging may eliminate certain log messages that could prove useful in a postmortem analysis of problem situations. The logging configuration can be tailored, as needed, to satisfy site or shift requirements.

Each SSM graphical user session (hpssgui) and command line session (hpssadm) can write a session log file of the SSM error and informational messages issued during the session. These are not the

same as the alarm, event, and other HPSS log messages described above and are not described in this chapter. See the **-S** option in the **hpssgui** and **hpssadm** man pages for details.

Log messages may be deposited in three places. Certain high-priority messages are transmitted to SSM and displayed on the *Alarms and Events* window. This allows the administrator to see events unfolding in close to real time. Log messages are written to the syslog subsystem. The administrator may filter, split, or configure the behavior of the HPSS logs using their system's syslog. HPSS log messages are logged using the "HPSS" identifier. Additionally, if the HPSSLOG_CONSOLE environment variable is enabled, log messages are also reflected to stderr.

Logging is a component which is utilized by HPSS servers and utilities. The logging component reads its log policy and then filters messages based on the configured log policies. Messages which pass the filter are logged to syslog, and may also be forwarded to the *Alarms and Events* SSM window.

For useful logging to take place, it is important that log policies and syslog are all configured and managed properly.

9.2. Log policies

For any HPSS server or utility which uses the HPSS logging facility, a log policy can be configured to specify which log message types will be sent to syslog.



If no log policy is configured for a server, the server will use the default log policy.

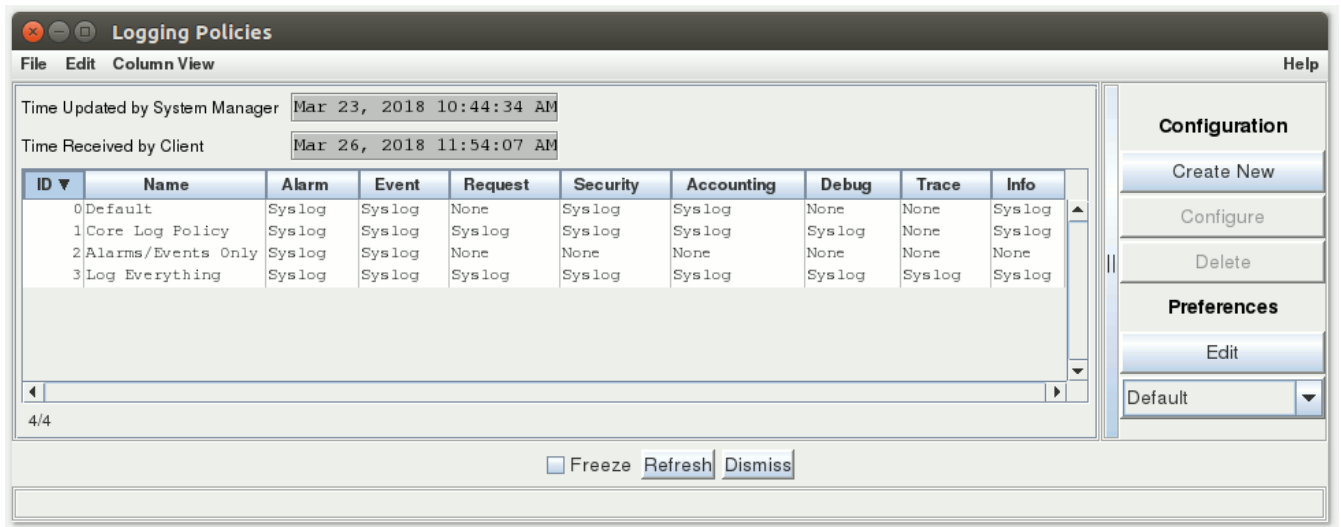
After changing any log policy information, including the Default Log Policy, changes will be reflected the next time the logging component rereads the logging policy. Processes using the HPSS logging component automatically reread their logging policy every thirty seconds.

9.2.1. Creating a log policy

A new log policy can be created and managed using the *Logging Policy* window (see *Section 9.2.2.1, "Logging Policy configuration window"*). This *Logging Policy* window can be accessed from the *Logging Policies* window (see *Section 9.2.2, "Logging Policies window"*). From this window, a log policy can be created with any unique name which is typed into the **Descriptive Name** field on the *Logging Policy* window.

Newly-created servers will use the default log policy. They can be assigned a different log policy in the *Section 5.2, "Server configuration"* window.

9.2.2. Logging Policies window



This window is used to manage all the log policies in the HPSS system.

To create a new log policy, click on the **Create New** button. To configure an existing log policy, highlight the policy in the list and click on the **Configure** button. When creating or configuring a log policy, the *Logging Policy* window will appear. To delete an existing log policy, highlight the policy in the list and click on the **Delete** button. If any servers still use the log policy, you will not be able to delete it. Note that the default log policy cannot be deleted and the **Delete** button is grayed out when selecting that server.

Field descriptions.

Logging Policy List

Name

The descriptive name of the log policy.

Log Policy List table columns

For details on the fields listed for each log policy, refer to the *Logging Policy* window (see *Section 9.2.2.1, "Logging Policy configuration window"*).

Configuration buttons.

Create New

Opens a *Logging Policy* window containing default values for a new policy.

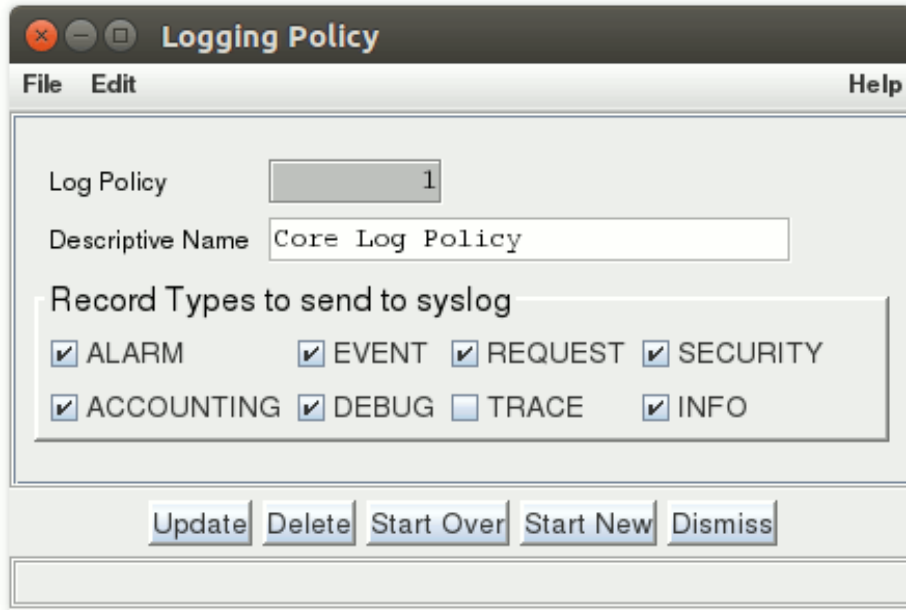
Configure

Opens the selected log policies for configuration editing.

Delete

Deletes the selected log policies.

9.2.2.1. Logging Policy configuration window



The *Logging Policy* window is used to manage a log policy.

When creating a new log policy, the **Descriptive Name** field will be set to the form "Log Policy #", where # matches the ID number for the new policy. A set of record types to send to syslog matching the default log policy will be selected.

For logging of a record type to be enabled, the check box located to the left of the log record type must be selected. The log state for each record type can be toggled by clicking on the check box.

After creating or modifying a logging policy, including the default, it may take up to thirty seconds for changes to propagate to the entire system. For example, if you modified log policy A, any servers using that policy would begin logging according to that policy within thirty seconds.

Field descriptions.

Log Policy

Log policy identifier

Descriptive Name

The descriptive name of the HPSS server or utility to which this log policy applies.

Record Types to send to syslog

Record types that are to be sent to syslog for the specified server. If a check box is selected for a record type, that type will be sent to syslog. The record types are:

ALARM

A message reporting a situation that requires administrator intervention. It is recommended that this always be selected.

EVENT

An informational message (for example, subsystem initializing or subsystem terminating) about a significant occurrence in the system that is usually not an error. It is recommended that this always be selected.

REQUEST

A message which reports the beginning and ending of processing of a client request in a server. It is recommended that this type not be selected except for short periods as an aid to isolating a problem.

SECURITY

A message reporting security-related events (for example, authorization failures). It is recommended that this always be selected.

ACCOUNTING

A message which logs accounting information.

DEBUG

A message reporting an internal system event that can aid in troubleshooting problems. For troubleshooting, these messages are important record types because they provide more detailed information about the root cause of an error. It is recommended that this always be selected.

TRACE

A message reporting procedure entry and exit events. It is recommended that this type not be selected except for short periods as an aid to isolating a problem.

INFO

A message reporting progress and status for more common, longer-running requests, as opposed to certain less frequent, more significant operations covered by EVENT log messages.



It is recommended that at least the Alarm, Event, Security and Info record types be selected for all servers while they are running normally. Core Servers and Movers should also log Debug messages, and additional record types for debugging purposes should be selected for servers experiencing problems. If system performance is being degraded by excessive logging, first turn off Trace, Debug and Request record types for all servers. Other record types can be turned off as necessary; however, the HPSS administrator may lose useful debugging information when problems occur. To facilitate diagnosis of data transfer problems, it is recommended that Debug logging always be turned on for all Movers.

9.2.3. Changing a log policy

A server's log policy can be modified to control the volume of messages to the chosen logging destinations. Typically, during normal operations, the level of logging may be decreased to only Alarm, Event, and Security to reduce overhead; however, while tracking an HPSS problem, it may be desirable to include additional log message types such as Debug, Request and Trace to obtain more information. As stated previously, it is recommended that Trace not be selected except for short periods as an aid to isolating a problem.

After a log policy has been modified, any servers using it will have their logging behavior changed within thirty seconds.

9.2.4. Deleting a log policy

There are two ways a log policy can be deleted:

- Highlight the log policy in the *Logging Policies* window (see Section 9.2.2, “*Logging Policies window*”) and click the **Delete** button.
- Open the *Logging Policy* window (see Section 9.2.2.1, “*Logging Policy configuration window*”) for the appropriate log policy and click the **Delete** button.

You will not be permitted to delete a log policy if any servers are configured to use it. Also, the default log policy cannot be deleted.

After deleting a log policy, any changes to logging, for example, reverting to use the default log policy, will begin within thirty seconds.

9.3. Managing syslog

This section describes configuring and viewing HPSS system logs. Note that for servers and processes using the logging subsystem, logs will appear in a location defined by that system’s syslog configuration. HPSS log messages are text-based and may be manipulated and viewed with standard UNIX tools.

9.3.1. Configuring syslog

This documentation will assume rsyslog is being used by default. Notes and information for other syslog implementations may be found inline where there are interesting differences between implementations.



Configuring logging requires careful consideration of the site’s needs and resources. Too much logging can exhaust disk space, hide relevant information, and degrade performance. Too little logging won’t provide enough information for resolving system problems when they arise. Use the options in your logging policies to adjust the volume of logging information recorded in a way that is commensurate with your site’s needs and resources.

HPSS log messages are sent to syslog using the "HPSS" identifier. Any file which contains HPSS system logs is referred to as an "HPSS log file". The syslog configuration has full control over the file or files that HPSS log messages appear in, and system utilities such as **logrotate** control the log retention and log archival strategy.



When using **vim** to view HPSS log files, `:set syntax=messages` will turn on syntax highlighting similar to `/var/log/messages`.

9.3.2. Recommendations for syslog configuration

The rsyslog service has a wide variety of configuration options which are largely outside the scope of this document; however, a few items may be useful in initially setting up syslog.

With no additional configuration, HPSS logs will go into the default syslog location, which is likely `/var/log/messages`. By default, SSM will look for logs in `/var/hpss/log/HPSS.log`. Mirroring HPSS alarms and events into `/var/log/messages` may be desirable.

Given the relatively verbose nature of the HPSS logs, it may be better for the majority of the HPSS logs to go elsewhere. The recommended syslog configuration for the HPSS Core Server is located in: `$HPSS_ROOT/config/templates/system_files/hpss.conf.rsyslog.template`

With the template configuration, the Core Server will act as a syslog server where all HPSS messages will go (in `/var/hpss/log/HPSS.log`). Alarms and events will also go to `/var/log/messages`. Remember that HPSS filters logging by server based upon that server's log policy. The default port is 514.

Remote systems must also configure a syslog configuration. This causes the remote system to log back to the HPSS Core Server. The remote template needs to be updated to include your HPSS Core Server in the forwarding line. The recommended syslog configuration for the HPSS Core Server is located in: `$HPSS_ROOT/config/templates/system_files/hpss.remote.conf.rsyslog.template`

HPSS rsyslog and logrotate templates can be automatically installed using the `$HPSS_ROOT/config/configure_logging` tool. This will copy any necessary template files into the proper locations for rsyslog and logrotate.



Note if your site uses syslog-ng, you should not use `configure_logging`. It is important to review the logging and log rotation configuration files before beginning use in a production environment.

9.3.3. Interactions between HPSS and syslog

The default configuration for *rsyslog*, the version of the syslog system that ships with RHEL systems, "escapes" nonprintable control characters in its input by replacing them in its output stream with their octal representation preceded by a hash symbol ("#"). For example, TABs are replaced by "#011" and newlines are replaced by "#012".

If this behavior is not desired, the rsyslog configuration can be changed to convert such characters to spaces instead. To get this effect, update the rsyslog configuration with the following lines:

```
$EscapeControlCharactersOnReceive off
$template HPSS_Format,"%TIMESTAMP% %HOSTNAME %msg:::space-cc%\n"
```

Note that after changing the rsyslog configuration, it will be necessary to "kill -HUP <rsyslog PID>" or restart the rsyslog server.

9.3.4. Configure syslog message size



To avoid message truncation, the maximum message size should be increased from the default.

HPSS log messages, especially IOD/IOR logs, can be extremely large (256 KB). To avoid log message truncation, it is necessary to increase the syslog maximum message size. It is expected that a syslog maximum message size of 256 KB will support nearly all of the log messages HPSS is likely to generate for the foreseeable future without truncation.

```
/etc/rsyslog.d/hpss.conf:
```

```
...
$MaxMessageSize 256k
...
```

```
/etc/syslog-ng/syslog-ng.conf:
```

```
options
{
...
    log_msg_size(262144);
...
};
```

9.3.5. Configure multi-part messages

Messages that happen to be larger than the `MaxMessageSize` are simply truncated in syslog. HPSS can work around this for extremely large messages by dividing the message into multiple syslog calls. The limitations on this are configured with the environment variable `HPSS_LOG_MSG_LEN` which can be set in `/var/hpss/etc/env.conf`.

HPSS will log any messages larger than `HPSS_LOG_MSG_LEN` in parts. Each part will have the same header content with the exception of the syslog timestamp. The message payload (`#msg=`) will begin with `MULTIMSG` followed by a unique identifier and the part number. The remainder of the payload contains part of the message that is being logged. The final part is noted by an exclamation point (!) appended to the `MULTIMSG` tag (`MULTIMSG!`). Below is an example of a three-part message.

```
MULTIMSG 93e0011a34bc234b958e00e3376353b3 0 <message content>
MULTIMSG 93e0011a34bc234b958e00e3376353b3 1 <message content>
MULTIMSG! 93e0011a34bc234b958e00e3376353b3 2 <message content>
```

The default `HPSS_LOG_MSG_LEN` is 260000. This is slightly less than the default message size in the HPSS syslog template. When setting `HPSS_LOG_MSG_LEN`, be aware that this does not include the timestamp and hostname prepended by the syslog call. Thus, `HPSS_LOG_MSG_LEN` should never be set to match the message size limit on your syslog implementation. It should be slightly less to account for this additional space in the messages.

Multi-part messages will appear in the same locations as expected (e.g. SSM), but it currently does not recombine them into the original message. Messages that are not divided into parts will be unaffected. Typically, only IOD and IOR messages are large enough to hit these limits. The `hpss_log_formatter` tool has been updated to reconstruct messages as it reads the log file. It will print the message once it has all of the parts or when the end of file is reached. Incomplete messages will be marked as such.

9.3.6. Restart Syslog After Configuration Change

Note that after changing the syslog configuration, it will be necessary to restart the syslog service using **service** or **systemctl**. On Linux, using `syslog-ng` or `rsyslog`, for example:

```
sudo service rsyslog restart
```

or

```
sudo service syslog-ng restart
```

9.3.7. Managing rate-limiting and duplication

Syslog and Linux in general supports several rate-limiting strategies including dropping messages and collapsing duplicate log messages. These could potentially impact HPSS. Generally for HPSS logs, we would like to eliminate rate-limiting on the Core Server and Mover systems.

The following, when added to `/etc/rsyslog.conf`, will disable rate-limiting for that server.

```
$SystemLogRateLimitInterval 0
$SystemLogRateLimitBurst 0
```

In addition, sometimes the **journal** can rate limit messages. Those settings can be modified in `/etc/systemd/journald.conf`:

```
RateLimitInterval=0
RateLimitBurst=0
```



It is possible to configure the interval and burst to match your system needs; however, this needs to be configured at the high end of the expected volume, or else you may begin losing messages. Check your logs during periods of high activity to ensure that messages are not being dropped.

If messages are being dropped, you may begin to see messages such as:

```
imjournal: begin to drop messages due to rate-limiting
Suppressed 1269 messages from /user.slice/user-648.slice
```

If you began to see these messages, adjust your rate limiting in the appropriate configuration files and restart syslog.

In order to disable message reduction via duplication, you can set the following in the HPSS syslog configuration file:

```
$RepeatedMsgReduction off
```

9.3.8. Considerations For SELinux

Systems which are protected by SELinux require additional configuration so to make the HPSS log file accessible to syslog. Otherwise, SELinux will block rsyslog from appending to the HPSS log file.

This can be resolved by giving rsyslog (or your local syslog equivalent) access to the appropriate syslog file targets. For example, if logging to `HPSS.log`, then run:

```
# chcon -t var_log_t HPSS.log
```

9.3.9. Managing log archiving

HPSS takes advantage of the Linux service **logrotate** to facilitate the rotation of its log files. Other Linux services could be used to provide this feature; however, this section will focus on a recommended configuration for **logrotate**, which will provide log rotation and archival services for HPSS.

As logs are archived via **logrotate**, they will be moved into HPSS with a name based upon the HPSS log times that appear in the log. The log file names will appear in one of the following forms:

Table 9.1. Log file formats

File Name Format	Location	File Status
logfile	File system	File is being filled
logfile.1	File system	File has been rotated once
logfile.1.gz	File system	File has been rotated once and compressed
logfile.YYYY.mmdd.HHMMSS.YYYY.mmdd.HHMMSS	HPSS Log Archive	File has been archived

Note that timestamp date ranges are appended to the log file only when it has been archived into HPSS.

The logrotate policy can set the number of times to rotate log files before deleting them for convenience. When log archiving is included in the logrotate policy, old log files can be retrieved from the archive.

If log archiving is not included in the logrotate policy, then the logrotate policy indicates the number of files to be retained on disk. Once a log file is deleted, it cannot be retrieved.

The administrator needs to consider whether archiving of the log files is needed, what the maximum size for log files should be, how many log files should be retained on disk, and then appropriately set a logrotate policy. See **man logrotate** for details on **logrotate**. Note that if log archiving fails for an HPSS log file, additional log files cannot be archived until that log file is archived.

If **logrotate** is not set up to archive HPSS log files, and then later the administrator decides to enable log archiving, the administrator should first manually archive all HPSS log files beyond the first rotation (for example, `HPSS.log.2` and beyond). Then they may enable log archiving via **logrotate** without losing any existing log information. See **man hpss_log_archive** for usage and options.

If **logrotate** is set up to archive HPSS log files and later log archiving is disabled, then no further HPSS logs will be archived.

The directory to be used for logging in the HPSS archive must be writable by the **hpss_log_archive** user so that the log archive tool will be able to create log files and new year, month, and day directories as needed. In order for this to occur, both the POSIX-style permissions on the directory must allow write permission and, if account validation is enabled, that user must be in the ACCTVAL table with a default account code. Use the **hpss_avaudit** utility to check this configuration and change, if necessary.



*If log archiving is set up via **logrotate**, but log file archiving has temporarily stopped working for some reason (for example, due to issues writing to the archive or parts*

*of HPSS being temporarily down), HPSS log files will not rotate and will continue to accumulate until log file archiving begins working again. Care should be taken when temporarily disabling log archiving to allow for log rotation to occur as this could result in the loss of log records. Any logs which are rotated without being archived via **logrotate** can be archived using the **hpss_log_archive** tool manually. See the `hpss_log_archive.7` man page for usage.*

9.3.10. Recommendations for logrotate configuration

It is recommended to use the **logrotate** utility for managing the rotation and archiving of HPSS log files. Each HPSS log file should have a logrotate definition in `/etc/logrotate.d/hpss`. See **man logrotate** for information on **logrotate** options. The template file for the recommended logrotate configuration is located in `$HPSS_ROOT/config/templates/system_files/hpss.logrotated.conf.template`.

The template logrotate definition has the following features:

- It will keep five log files locally before deleting them.
- Logs will become eligible for rotation after they are 100 MB in size.
- Logs will be owned by root and have 0600 permissions by default.

There are two script stanzas: prerotate and postrotate. Let's start with postrotate first.

```
postrotate
killall -HUP rsyslogd
/opt/hpss/bin/hpss_log_archive -f $1.1
endscript
```

Every line after postrotate is a script, which will be run after the log has been rotated (renamed). The first thing we do is SIGHUP rsyslog to stop logging to the rotated file and force it to open the new file.

Now for prerotate:

```
prerotate
/opt/hpss/bin/hpss_log_archive -R -f $1.1
endscript
```

Prerotate handles the situation where a prior postrotate operation failed. The `-R` (for recovery) parameter tells **hpss_log_archive** that if the file doesn't exist, just return success. If there is a `logfile.1` out there waiting to be archived though, it will attempt to archive it. If that is successful, then we will log a message indicating we're continuing on with archiving the current log. Note that the call to `hpss_log_archive` must be the final step in the prerotate and postrotate scripts, or their exit codes retained, so that if the archive step fails, that the active log will continue on *without rotating*.



The **hpss_log_archive** tool has additional options to control the way that files are written into HPSS. See the **hpss_log_archive** man page for details.

HPSS rsyslog and logrotate templates can be automatically installed using the `$HPSS_ROOT/config/configure_logging` tool. This will copy any necessary template files into the proper locations for rsyslog and logrotate.



In order to properly handle log archive failures, the call to **hpss_log_archive** must be the final call in each script, or must be wrapped in a script which retains and returns any errors from **hpss_log_archive** - the exit code of the final script directive determines logrotate's behavior for prerotate and postrotate, any intermediate errors are ignored.



Note if your site uses syslog-ng, you should not use `configure_logging`. It is important to review the logging and log rotation configuration files before beginning use in a production environment.

A cron job should be configured to run **logrotate**. You may want to have it run every five minutes against the HPSS definition to check the logs, for example. If the log file doesn't match the requirements laid out in the definition, then **logrotate** won't do anything.

```
*/5 * * * * /usr/sbin/logrotate /etc/logrotate.d/hpss
```



The **hpss_log_archive** tool currently does not work with the `dateext` logrotate option. The `nodateext` option must be used for the **hpss_log_archive** tool to process rotated log files.

9.3.11. Viewing HPSS logs

HPSS provides the ability to retrieve and examine HPSS log records as a means of analyzing the activity and behavior of HPSS.

To effectively view log records, the user should know what kinds of activities or behaviors are of interest. Since it is possible to configure log policies to include almost all activity in HPSS, in most cases it will be easier to examine and analyze logs by filtering out any records that are not of interest. Since the log files are stored in text format, this can be done using standard UNIX text processing tools.

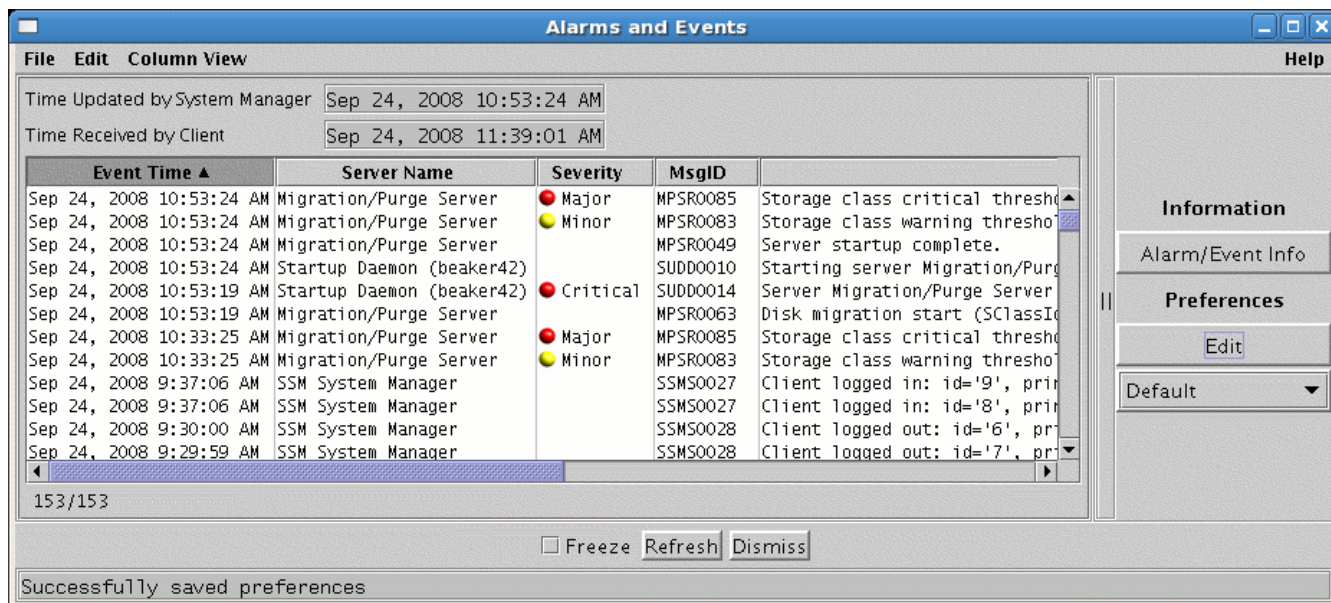
When the HPSS log file, referred to here as `logfile`, reaches the maximum size indicated by its logrotate policy, **logrotate** will rotate all already rotated logs (`logfile.1-logfile.N`) such that `logfile.1` becomes `logfile.2` and so on, and anything beyond the logrotate configured rotate count is deleted. Additionally, **logrotate** will rename `logfile` to `logfile.1`. Note that until syslog has been restarted or signaled with `SIGHUP`, it will keep writing to `logfile.1`. After syslog has switched log files, if log archiving is configured in the logrotate policy, **logrotate** calls **hpss_log_archive** with the file that was just rotated (`logfile.1`) and **hpss_log_archive** writes that file into HPSS with its start and end timestamps.

If the log files of interest have rolled off the disk, as long as archiving is configured, they can be retrieved from a date-based directory structure in HPSS. For example, `/log/yyyy/mm/dd`.

9.4. Managing SSM alarms and events

This section describes how to manage SSM alarms and events.

9.4.1. Alarms and Events window



This window displays a number of the most recent alarm and event messages which have been received by SSM. It also allows you to view individual messages in greater detail by selecting the message and clicking the **Alarm/Event Info** button to bring up the *Alarm/Event Information* window.

Field descriptions.

This list displays a column for each field shown on the *Alarm/Event Information* window. See *Section 9.4.2, "Alarm/Event Information"* for the field descriptions.

9.4.2. Alarm/Event Information

The screenshot shows a window titled "Alarm/Event Information" with a menu bar containing "File", "Edit", and "Help". The window contains the following fields:

- ID: 49817
- Log Type: ALARM
- Severity: ● Major
- Event Time: Mar 29, 2018 9:54:45 AM
- Server Name: Migration/Purge Server
- Routine: mps_CheckSClassThreshold (line 1286)
- PID: 3672
- Node: devNode-3.clearlake.ibm.com
- Process: hpss_mps
- Request ID: 0
- Error Code: 0
- MsgID: MPSR0085
- Text: Storage class critical threshold exceeded (SClassID 4, SubSysID 1).

At the bottom of the window is a "Dismiss" button.

This window displays all the details of the alarm or event selected from the *Alarms and Events* window.

Field descriptions.

ID

A sequence number assigned to the log message by SSM. This ID is not used outside of SSM.

Log Type

General class of the message. May be either ALARM or EVENT.



Severity

For events, this field will be blank. For alarms, it will be one of the following:

- Warning
- Minor
- Major
- Critical

- Indeterminate
- Cleared

These may be accompanied by a color status indicator:

-  (Red) for a Critical or Major alarm
-  (Yellow) for a Minor or Warning alarm
- No indicator for events and other alarm types

Event Time

The date and time the message was generated.

Server Name

Descriptive name of the HPSS server or utility that logged the message.

Routine

Name of the function that was executing when the message was logged.

PID

Process ID of the process that logged the message.

Node

Name of the host where the message was generated.

Process

Name of the process that generated the message.

Request ID

Request ID generated by the process logging the message.

Error Code

The error code associated with the problem underlying the message.

MsgID

The eight-character ID code for this message, consisting of a four-character mnemonic identifying the type of server or subsystem which issued the message followed by a four-digit message number. (The message number is not the same as the **Error Code** field.) Message IDs may be used to look up messages in the *HPSS Error Manual*.

Text

The text of the alarm or event message. Only up to 240 characters are displayed. Look at the `HPSS.log` file to view the entire message.

9.4.3. Diagnosing HPSS problems with alarms and events

Events displayed on the *Alarms and Events* window provide an indication that a significant event has occurred in a server and the event may be of interest to the administrator. An alarm, however,

indicates that a server has detected an abnormal condition. The user should investigate the problem as soon as possible to ensure timely resolution. The user may use the information provided by the alarm to obtain further information on the problem as follows:

- Use the alarm message number to look up the alarm information in the *HPSS Error Manual*. For each documented message, the manual provides more detailed information on the problem and its possible source. The manual also provides recommendations on how to resolve the problem.
- If additional information is needed, use the alarm timestamp to identify the log files containing the log messages received around the time the problem was reported. In addition to obtaining the messages logged by the server in question, it may be necessary to obtain log messages from other HPSS servers that interface with the server. Refer to *Section 9.3.11, “Viewing HPSS logs”* for more information on locating the HPSS log messages.

9.4.4. Controlling SSM log message handling

This section describes the options available to control SSM handling of HPSS log messages.

This does not include the SSM session log. See the **-S** option on the **hpssgui** or **hpssadm** man page for a discussion of SSM session logging.

9.4.4.1. Controlling the System Manager log message cache

By default, the SSM System Manager stores the alarm and event messages it receives in an internal memory cache. This cache can be configured instead to be kept in a disk file by defining the environment variable `HPSS_SSM_ALARMS` with the desired name of the cache file. The default for `HPSS_SSM_ALARMS` is defined in `hpss_env_defs.h` as `NULL`. SSM will revert to the internal memory cache if it cannot access the specified cache file for any reason.

The site may set the `HPSS_SSM_ALARMS` environment variable to any UNIX file that has read/write access for user *root* on the machine where the SM is to be run (since the SM runs as user *root*). A good path might be `/var/hpss/log/alarm_buffer`, for example.

Using a disk cache makes SSM alarms and events "persistent"; that is, if the System Manager is restarted, SSM will be able to display all old messages still in the disk cache. Caching the alarms and events in memory means the alarm cache starts out empty each time the System Manager is restarted.

By default, 2000 messages are cached before the oldest are discarded. This value can be modified by setting the `HPSS_SSM_ALARMS_DISPLAY` environment variable to the desired number of messages to be retained and restarting the System Manager.

Once an old message is discarded, it cannot be accessed anymore from the *Alarms and Events* window. Older messages can be accessed by examining the `HPSS.log` file, a log rotated `HPSS.log` file, or by retrieving even older log files from the HPSS archive.

Since the messages are cached and managed by the System Manager, all SSM users connected to the System Manager will see the same messages; however, preferences selected on the *Alarm and Event Preferences* window are applied on a user-by-user basis.

Messages can be delivered to the System Manager and stored in the cache only while the SSM System Manager is running. Any messages generated while the System Manager is down will not be

accessible from the HPSS *Alarms and Events* window (but they should be accessible in the local log, central log, or archived log files as long as the logging service processes were running when they were generated). This means that there can be gaps in the list displayed in the *Alarms and Events* window if the System Manager is shut down and restarted while the rest of HPSS is up.

9.4.4.2. Controlling log messages displayed by **hpssgui** and **hpssadm**

The **hpssgui** and **hpssadm** programs each keep an internal cached copy of the alarm and event list. The list is maintained by regular polling requests to the System Manager. Review and reference the **hpssgui** and **hpssadm** man pages for further details on the settings discussed in this section.

The polling interval is set by the HPSSSSM_UI_ALARM_RATE environment variable, which may be overridden by the **-A** option to the **hpssgui** or **hpssadm** startup script. The default is to poll the System Manager every five seconds for new log messages.

Each polling request returns only "new" messages, those which have not already been retrieved by the **hpssgui** or **hpssadm** in a previous polling request. If there are no new messages in the System Manager log cache, no messages are returned.

By default, the maximum number of messages retrieved by each polling request is 2000. If there are more than 2000 new messages in the System Manager log message cache, only the 2000 newest will be returned. This means that if new messages are arriving rapidly at the System Manager and the **hpssgui** or **hpssadm** polling interval and maximum number of messages values are not set appropriately, there can be gaps in the **hpssgui** or **hpssadm** alarm and event list.

The value for the maximum number of messages to return by each polling request may be redefined by the HPSS_SSM_ALARM_GET environment variable or by the **-G** option to the **hpssgui** or **hpssadm** startup script. The command line option takes precedence over the environment variable. This option could be used, for example, to reduce the size of each data request on a slow network.

The internal cached alarm and event list is displayed by the **hpssadm** program by means of its "alarm list" command. This command has a **-c** option to specify how many of the most recent log messages in the internal copy to display. If more messages are requested than exist in the internal list, the full internal list is displayed. See the **hpssadm** man page for details.

The internal cached alarm and event list is displayed by the **hpssgui** in the *Alarm and Events* window as described in *Section 9.4.1, "Alarms and Events window"*. The maximum number of log messages displayed in the window is set by the HPSS_SSM_ALARM_DISPLAY environment variable, which may be overridden by the **-N** option to the **hpssgui** startup script. The default is to display at most 2000 messages in the list.

If the maximum number of messages to retrieve (HPSS_SSM_ALARM_GET or **-G** option) is greater than the maximum number of messages to display (HPSS_SSM_ALARM_DISPLAY or **-N** option), it will be reset to the same value as the maximum number of messages to display. This is true for both the **hpssgui** and the **hpssadm**, although the **hpssadm** does not use the maximum number of messages to display value for anything else.

If new messages are arriving regularly, the messages displayed by the **hpssgui** *Alarms and Events* window will constantly be moving downward at set intervals. This can make it difficult to select

a specific message if the window is updating frequently. This problem can be minimized by being careful to click on a message immediately after an update cycle, or it can be eliminated entirely by freezing the display by checking the Freeze check box at the bottom of the *Alarms and Events* window.

If many messages are arriving very rapidly, or if you leave the Freeze button on, it is possible for the display to become stale. A stale message is one which is still visible, but which has been discarded from the cache. If you click on such a message, an error box will appear, telling you that the selected message is no longer available.

Messages are displayed in the order in which they are received by the System Manager. They are not sorted by the event time displayed with each message, so if many messages are generated in a short period of time, it is possible to see messages which are displayed slightly out of chronological order.

Each **hpssgui** user can filter and sort what is shown on his own *Alarms and Events* window independently of any other **hpssgui** user. See *Section 3.10, “SSM list preferences”* for information on filtering and sorting alarms. Filtering and sorting of the *Alarms and Events* list are not yet available for **hpssadm** users.

9.4.5. Media Access Logging

HPSS logs individual media access for I/O operations. This feature is always on, but the logging only occurs if the INFO log is enabled for the Core Server.

This feature allows an administrator to understand which files may have been written to or read from a particular disk or tape volume, which tape drives may have been used to read or write a file, and when those operations occurred. This can assist in certain data investigations, such as discovering files which could have been impacted by bad firmware.

Fields in the start access log:

Media Access

Media access logs contain an initial tag to differentiate them for filtering or redirection.

Operation

The operation that is beginning; for example, read, write, stage, copy, or migration.

BFID

The Bitfile ID for the file being logged.

Fileset

The fileset name for the file being logged.

Path

The path to the file from the fileset root.

Fields in the end access log:

Media Access

Media access logs contain an initial tag to differentiate them for filtering or redirection.

Operation

The operation that is completing; for example, read, write, stage, copy, or migration.

Media Read

In the format <Type><DeviceId>:<Volume>. Indicates the type (T for tape, D for disk) of media, the device ID that was used, and the one or more volumes read. If multiple volumes were read they will be separated by a comma.

Media Write

In the format <Type><DeviceId>:<Volume>. Indicates the type (T for tape, D for disk) of media, the device ID that was used, and the one or more volumes written. If multiple volumes were written they will be separated by a comma.

BFID

The Bitfile ID for the file being logged.

Fileset

The fileset name for the file being logged.

Path

The path to the file from the fileset root.

Result

Indicates the result as a string along with the numeric error code.

Media Access logs look like:

```
May 30 10:02:07.793192 hpss_core(HPSS)[756]:: #msgtype=INFO
#server=Core Server@server #func=bfs_StageThread
(line 1478) #rc=0 #msgid=CORE2188 #reqid=6068557350692913156
#msg=Media Access : Stage begin : BFID
x'00000539070000000801E841C3C0A743B01E34' : FilesetRoot.29 :
./path/to/file
May 30 10:02:07.793801 hpss_core(HPSS)[756]:: #msgtype=INFO
#server=Core Server@server #func=bfs_StageCleanup
(line 3046) #rc=0 #msgid=CORE2189 #reqid=6068557350692913156
#msg=Media Access : Stage end : read {T4:JD021700} : wrote
{D2:MY000200,D1:MY000100} : BFID
x'00000539070000000801E841C3C0A743B01E34' : FilesetRoot.29 :
./path/to/file : Result No error (0)
```

Media access logs exist for file read, write, migrate, batch migrate, stage, async stage, and copy operations. Media access logging is currently unavailable for Repack.

9.5. Log files of various types - a comprehensive list

HPSS generates log files of its own; however, a number of log files of various types are generated in the operation of an HPSS system. There are also files that are not log files, strictly speaking, but which grow over time and have the potential to fill the disk where they reside. This section attempts to provide a comprehensive list of such files along with guidance for managing them.

/var/hpss/adm/hpssd.failed_server:

A record is written to this file each time an HPSS server exits abnormally. Normally, the growth of this file will be very slow, but eventually it could fill the disk, so it should be monitored and removed if it threatens to become too large.

/var/hpss/adm/core:

If an HPSS server suffers a severe internal error and crashes, a diagnostic core dump will be placed in `/var/hpss/adm/core/<Server-Name>/core.<pid>`. These core files can be very large and if enough of them accumulate, the disk could be filled. This area should be checked periodically and cleaned out if necessary.

/var/hpss/ftp/adm/hpss_ftpd.log:

PFTP uses `/var/hpss/ftp/adm/hpss_ftpd.log` for Client API debugging. Debugging is enabled by setting

```
Client API Verbose Value = <n>
```

in `HPSS.conf`, where `<n>` is the debugging level desired. The file input is provided directly by the Client API. If debugging is left on, this file can also grow without bounds, so it should only be used for debugging and removed when debugging is completed.

/var/hpss/hpssdb/sqllib/db2dump/db2diag.log:

DB2 writes diagnostic messages to `db2diag.log`. If this file gets too large, it can be renamed and optionally compressed for archival. Once a copy has been archived, the renamed file should be removed from the system.

To manage the size of this file, you can periodically archive and remove it or use a log rotation tool of your choice (for example, `/usr/sbin/logrotate`).

/var/hpss/log/krb5kdc.log:

Kerberos information and errors are written to this file. Its location is defined in the Kerberos configuration, usually in `/etc/krb5.conf`. This file will only be generated if you use Kerberos. To manage its size, you can periodically archive and remove it or use a log rotation tool of your choice (for example, `/usr/sbin/logrotate`).

/var/hpss/mps/<report-files>:

The Migration/Purge Server will write daily reports in `/var/hpss/mps` if the MPS Report File Name field in the Migration/Purge Server Specific Configuration is not empty. Unless this directory is monitored and old reports are removed, these report files have the potential of filling the disk.

/var/hpss/tmp/gasapi.diag:

See the description of `secapi.diag` below. This file behaves exactly the same way.

/var/hpss/tmp/secapi.diag:

Normally, on HPSS systems, this file does not exist. It can be used to troubleshoot security issues in HPSS if needed. To use it, issue the command

```
touch /var/hpss/tmp/secapi.diag
```

Once the file exists, security logs will be appended to it as they arise. The file contents can be viewed with any text file viewer.

A command like the following will truncate the file:

```
/bin/echo -n "" > /var/hpss/tmp/secapi.diag
```

If this file exists and HPSS is running, the file should not be deleted. Deleting a file while one or more processes have it open will remove the file's directory entry but leave the file's data segments on disk still allocated and occupied. As the running processes continue writing to it, the file will grow invisibly, eventually consuming all available disk space.

If this should happen, to recover, it is necessary to shut down all HPSS processes. Once all the processes holding the phantom file open terminate, the operating system will notice that the disk blocks are associated with a file that no longer exists and recover them.

/var/log/wtmp, /var/adm/wtmp:

UNIX writes a record to this file any time a user logs in to or out of the system. These records can be viewed with the **last(1)** utility. If this file grows too large and needs to be removed, with the appropriate privilege, it can be truncated by redirecting the output of a silent command (that is, one that produces no output):

```
cat /dev/null > /var/log/wtmp
```

To manage the size of this file, you can periodically archive and remove it or use a log rotation tool of your choice (for example, `/usr/sbin/logrotate`).

/var/hpss/tmp/scsi.diag:

`device_scan`, the SCSI PVR, the Mover, and other tools that interact with the HPSS SCSI layer may log SCSI diagnostic messages to this log file. To enable this logging, touch the `scsi.diag` file. This logging can be helpful in troubleshooting issues in the SCSI or device layer. This file will grow without bound. In order to cause all HPSS processes to stop logging to this file, touch `/var/hpss/tmp/noscsidiag`. HPSS processes logging to `scsi.diag` will stop logging within 30 seconds and close the file. This allows the `scsi.diag` file to be cleaned up or rotated easily. Removing the `/var/hpss/tmp/noscsidiag` file resumes all logging for any HPSS process that is logging SCSI diagnostics.

/var/hpss/tmp/<PID>.diag:

Several HPSS servers — including the Core Server, PVL, PVR, SSM, Mover, RAIT Engine, SD, and GK — will dump out internal runtime information to a log file named after the server instance's Process ID (PID). The server state dump information typically includes generic information such as RPC diagnostics and cached SQL queries. After those sections, the servers output custom listings of internal runtime information.

Sending a `SIGHUP` signal to one of the listed HPSS server processes will cause the state dump file to be written to `${HPSS_PATH_TMP}`. These log files are typically most useful during investigation and debugging. It can be particularly helpful to generate one of these state dump files while a problem is occurring and then provide the file to HPSS Support.

API debug log:

The Client API can generate debug logging to assist with application and configuration issues. For a full discussion of Client API logging, consult the *Client API Tutorials and Best Practices* chapter of the *HPSS Programmer's Reference*.

There are three ways to enable the API debug log:

- Add `HPSS_API_DEBUG=7` and `HPSS_API_DEBUG_PATH=/path/to/log` to `$HPSS_PATH_ETC/env.conf`.
- Export `HPSS_API_DEBUG=7` and `HPSS_API_DEBUG_PATH=/path/to/log` before running the program.
- Update the `hpss.api.resource` file in `$HPSS_PATH_TMP`. The format is:
7 /path/to/log

The logging begins for environment variable methods the next time a Client API instance is initialized. In the resource file case, the logging will not begin immediately, but the next time the API reads the resource file.

The logging stops for environment variable methods when the API instance ends. For resource logging, the resource file can be removed or the log level changed to "0" to disable logging. Logging will be disabled the next time the API reads the resource file. The resource file impacts *all* running instances on the system so care should be used before using this method. The log location can also be specified as `stdout`.

The API resource file is checked every 30 seconds.

Authentication files:

Some HPSS utilities backup certain authentication files before changing them. These files reside in `/var/hpss/etc` and are named `passwd`, `group`, and `shadow`. The backups have names of the form `<filename>.<epoch>` where `<epoch>` is a time stamp consisting of the number of seconds since the beginning of 1970. If these files see a lot of activity, it may be necessary to periodically remove the backup files to keep the disk from filling.

Mover Debug Logging:

To control which Movers generate debug output, what message types are printed, the format of the output, and where the debug output messages will be written you must create a `MoverDebug.conf` file in the `$HPSS_PATH_TMP` directory (the default is `/var/hpss/tmp/MoverDebug.conf`). This file should be created on the server on which the `hpss_mvr_adm` process is running. After this file is created, any Movers added to `MoverDebug.conf` will need to be restarted.

The format of the file is:

`'MOVER DESCRIPTIVE NAME' MASK FILE_PATH`

where:

- `MOVER DESCRIPTIVE NAME` is the descriptive name of the Mover for which to enable debugging. It *must* be enclosed in single quotes (').
- `MASK` is a hexadecimal string specifying the format and types of messages to output. The string is obtained by ORing the desired flags (found in `mvr_shmem.h`). Using `0x00FFFFFF` will enable all Mover logging options.
- `FILE_PATH` is either the pathname of a file that will contain the output or the string `stdout`, in which case the output will be written to stdout of the tty associated with the root shell that started the Mover or the SSM.

Any line starting with a `#` symbol is ignored.

Example:

```
'DiskMover1 (gauss)' 0x50027 /var/hpss/tmp/DML.out
'TapeMover (T10K)' 0x0F003F stdout
```

The Mover debug configuration file is read periodically (a few times a minute). Note that *only* the `MASK` value can be changed without restarting the Mover.

9.6. Disk Mover & VV I/O Metrics

The Core Server's Storage Service (SS) maintains I/O metrics against each of the disk movers and the disk VVs that those movers manage. The I/O metrics indicate how many read or write I/O tasks are ongoing against the movers and their disk VVs. As an I/O task begins, the SS increments the `OperationsInProgress` for the mover(s) and VV(s) involved in the I/O. When the I/O is complete, the SS decrements the `OperationsInProgress` for the mover(s) and VV(s) involved in the I/O. Note that read I/O and write I/O are weighted the same: the metrics are incremented and decremented exactly the same way regardless of whether reading from or writing to a mover and disk VV.

The SS uses these metrics to guide disk allocations. At any point in time, the number of `OperationsInProgress` for the various disk movers and VVs give a point-in-time snapshot of how busy the movers and disks are. This is a simple measurement of how many I/O tasks are currently in flight against the movers and their VVs. At point of allocation, the SS finds the mover and disk

VV with the lowest number of operations in progress and, if enough space is available, performs the allocation against that mover and disk VV. An allocation represents future I/O activity. Once the I/O for that allocation begins running, that's when the `OperationsInProgress` will be incremented.

The `TimeLastAllocated` is used as a tie breaker if movers and their VVs have the same number of operations in progress. In the case of a tie breaker, the least recently allocated mover and VV will be chosen. Again, note that read and write operations are given the same weight. If **mover A** has 8 `OperationsInProgress` (4 reads; 4 writes) and **mover B** has 8 `OperationsInProgress` (8 reads; 0 writes), they'll be considered to be equivalently busy.

Disk allocations can be guided to the least busy mover and disk VV. This is how disk write I/O is load balanced across the available disk movers and disk VVs. However, disk reads are not load-balanced. Those I/O tasks must occur on the movers and disk VVs that hold the requested data. The read I/O can't be shifted to other movers and VVs. So an influx of disk read requests can lead to an I/O imbalance across the disk mover complex. Note that both stages from disk and migrations from disk to tape are considered to be disk read I/O activity.

Metadata about the disk mover and VV I/O metrics — the most recent changes to the I/O metrics and the allocation decisions based upon them — can be viewed in the Core Server state dump `/var/hpss/tmp/<PID>.diag`. And a point-in-time snapshot of the I/O metrics is available via the unsupported `showdiskmaps` command line tool.

The I/O metrics metadata is ephemeral. It ages out pretty quickly. The metadata is interesting close in time to when the allocations and I/O occur, especially during various testing and debugging efforts. But after a short amount of time the metadata is no longer interesting. So, the in-memory metadata is regularly trimmed. This is an attempt to control how much memory is used by the I/O metrics metadata and to control how long the server state dumps are, especially on busy production systems.

The functionality consists of two related features:

- A log of all the allocations, reads, and writes and how they impact the metrics. This info is available in the Core Server state dump — obtained via sending `SIGHUP` to the `hpss_core` process. See annotated example snippet below Section 9.6.1, “Sample Log of Disk I/O Operations”.
- A point-in-time snapshot of disk and mover I/O metrics. This is available in the `showdiskmaps` tool's output. Below is an example for the 4-way VV `DLD21200`. Note that, as more concurrent I/O jobs run against HPSS, the `OpsInProgress` fields will increment.

```
VVID: Object: 000222e0-03-00000001-01e53bc3a4fcb612-35b1, Disk VV
      System ID: 140000, Server ID: 1
      DB2 format: x'000222E0030000000101E53BC3A4FCB61235B1'
PV Names: DLD21200 DLD21300 DLD21400 DLD21500
I/O Metrics
# VV
# # OpsInProgress:      4
# # TimeLastAllocated: 1629172518406154
# Mover ID 1041
# # OpsInProgress:      4
# # TimeLastAllocated: 1629172518406154
Number of Extents: 11, Map State: MAP_FREE, Storage Class ID: 140, Flags: None
Block Size: 67108864 bytes, Number of Clusters: 56079
Usable Length:      15053323501568 bytes
Actual Length:      15053659045888 bytes
Free Space:         15021111246848 bytes
```

```
Cluster Length:                268435456 bytes
Number of Map Blocks: 18
Number of slots in use: 13
Number of levels in use: 11
```

9.6.1. Sample Log of Disk I/O Operations

Here's a brief annotated snippet of the disk I/O operations log from a Core Server state dump taken soon after starting the Core Server. It shows some migrations (disk reads) and disk allocations.

```
===== Disk I/O Operations =====

# migration starts against mover 1041 & VV DLD21200

Start Disk I/O on Mover 1041

    08/16/2021 08:07:38 PM
    Request ID = e5eb0de439bd4043af33e5a790faac80
    VV = DLD21200
    Mover most recent allocation: 12/31/1969 04:00:00 PM
    VV most recent allocation: 12/31/1969 04:00:00 PM
    I/O Type = disk read
    I/O Operations in Progress:
        Prev Current
    Mover =      0      4
    VV =        0      4

# migration starts against mover 1029 & VV DLD20000

Start Disk I/O on Mover 1029

    08/16/2021 08:07:38 PM
    Request ID = 37175ef784bef741803e4738d56e3ac2
    VV = DLD20000
    Mover most recent allocation: 12/31/1969 04:00:00 PM
    VV most recent allocation: 12/31/1969 04:00:00 PM
    I/O Type = disk read
    I/O Operations in Progress:
        Prev Current
    Mover =      0      4
    VV =        0      4

# disk allocation against mover 1033 & VV DLD20400
# - a few notes:
#   o mover 1029 is skipped: it's "too busy" due to the above migration
#   o mover 1041 is skipped: it's "too busy" due to the above migration

Allocate Space on Mover 1033

    08/16/2021 08:08:29 PM
    Request ID = 29f655dcfc72004ab6827da6c558449f
    VV = DLD20400
    Mover most recent allocation: 08/16/2021 08:08:29 PM
    VV most recent allocation: 08/16/2021 08:08:29 PM
    I/O Type = disk write
    I/O Operations in Progress:
        Prev Current
    Mover =      -      0
```

```

V V =      -      0

Movers/VVs considered for allocation:
  Mover 1029:
    # Ops=4 LastAlloc=12/31/1969 04:00:00 PM
    # too busy
  Mover 1033:
    # Ops=0 LastAlloc=12/31/1969 04:00:00 PM
    # selected for allocation
    VV DLD20400:
      # Ops=0 LastAlloc=12/31/1969 04:00:00 PM
      # selected for allocation
  Mover 1037:
    # Ops=0 LastAlloc=12/31/1969 04:00:00 PM
    # used too recently
  Mover 1041:
    # Ops=4 LastAlloc=12/31/1969 04:00:00 PM
    # too busy

# disk allocation against mover 1037 & VV DLD20800
#   - a few notes:
#     o the same busy movers above are still listed as busy (the migrations are
#       still running)
#     o now mover 1033 is also "too busy" because of the previous allocation
#     o so, the remaining mover, 1037, is selected for the allocation

Allocate Space on Mover 1037

08/16/2021 08:08:29 PM
Request ID = 66fb03e3eefcab4c926ff025170445a0
VV = DLD20800
Mover most recent allocation: 08/16/2021 08:08:29 PM
  VV most recent allocation: 08/16/2021 08:08:29 PM
I/O Type = disk write
I/O Operations in Progress:
      Prev Current
Mover =      -      0
  VV =      -      0

Movers/VVs considered for allocation:
  Mover 1029:
    # Ops=4 LastAlloc=12/31/1969 04:00:00 PM
    # too busy
  Mover 1033:
    # Ops=0 LastAlloc=08/16/2021 08:08:29 PM
    # used too recently
  Mover 1037:
    # Ops=0 LastAlloc=12/31/1969 04:00:00 PM
    # selected for allocation
    VV DLD20800:
      # Ops=0 LastAlloc=12/31/1969 04:00:00 PM
      # selected for allocation
  Mover 1041:
    # Ops=4 LastAlloc=12/31/1969 04:00:00 PM
    # too busy

```

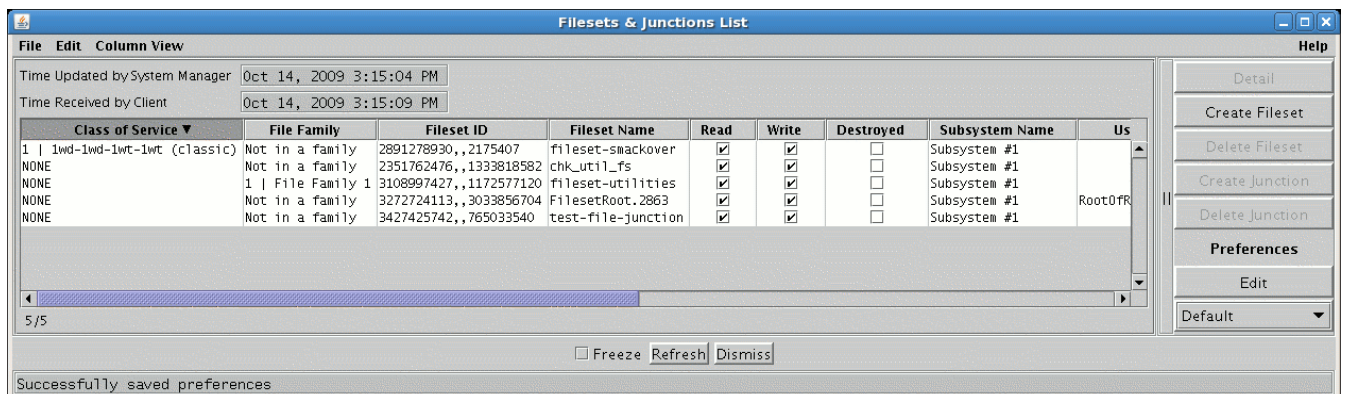
Chapter 10. Filesets and junctions

A fileset is a logical collection of files that can be managed as a single administrative unit, or more simply, a disjoint directory tree. A fileset has two identifiers: a human readable name, and a numeric fileset ID. Both identifiers are unique to a given HPSS realm. Filesets are often used for HPSS name space administration. For example they can be used if a subtree of the HPSS name space needs to be assigned to a particular file family or Class of Service.

A junction is a Core Server object, much like a symbolic link, that is used to point to a fileset. This fileset may belong to the same Core Server or to a different Core Server. When pointing to a different Core Server, junctions allow HPSS users to traverse to different subsystems.

When a Core Server initializes for the first time on an empty database, it creates the *root node* fileset. This is the fileset that contains the directory named "/" (forward slash). For many sites this may be the only fileset needed and no other filesets will ever be created.

10.1. Filesets & Junctions List



This window displays information about the filesets and junctions that are configured in the HPSS system. From this window, detailed information about existing filesets and junctions can be viewed, new filesets and junctions can be created, and existing filesets and junctions can be deleted.

To open the window, from the *HPSS Health and Status* window **Monitor** menu select **Filesets & Junctions**.

There may be multiple entries in the list for a given fileset meaning that there are multiple junctions that point to the same fileset. Each entry for a given fileset will have a unique Junction Name. If the Junction Name is blank, then there are no junctions that point to the particular fileset.

The **lsfilesets**, **lsjunctions** and **lshpss** utilities may also be used to list all the filesets and junctions in the system. For more information, refer to the man pages for each utility.

Field descriptions.

Filesets & Junctions List

Class of Service

The name of the Class of Service to which the fileset is assigned. If this field contains "NONE", the fileset has not been assigned to a Class of Service.

File Family

The name of the file family to which the fileset is assigned. If this field contains `Not in a family`, the fileset has not been assigned to a family.

Fileset ID

The ID number which identifies the fileset. A fileset ID is displayed as two double-comma-separated unsigned integer numbers.

Fileset Name

The unique name which has been assigned to the fileset.

Read

If checked, the fileset is available for reading.

Write

If checked, the fileset is available for writing.

Destroyed

If checked, the fileset cannot be written to or updated.

Subsystem Name

The subsystem name to which the fileset has been assigned. This is the subsystem of the Core Server which controls the fileset.

User Data

This field is available for storing up to 128 bytes of data. The information can be ASCII or binary.

Directories

The number of directories in this fileset.

Files

The number of files in this fileset.

Hard Links

The number of hard links in this fileset.

Junctions

The number of junctions in this fileset.

Sym Links

The number of symbolic links in this fileset.

Junction Name

The name of a junction that points to this fileset. This field may be blank if there are no junctions configured to point to this fileset.

Parent Fileset

The name of the fileset in which the junction resides. This field will be blank if the Junction Name is blank.

Buttons.**Detail**

Display the *Core Server Fileset Information* window for the selected filesets. This button is only enabled if one or more items are selected in the list.

Create Fileset

Display the *Create Fileset* window so that a new fileset may be created. This button is enabled all the time.

Delete Fileset

Delete the selected filesets. This button is only enabled if one or more items are selected in the list.

Create Junction

Create a junction to the selected fileset. This button is only enabled when there is exactly one item selected in the list. The fileset name and ID from the selected fileset will be used to fill in the information on the *Create Junction* window. The user will only have to supply the Junction Name on the *Create Junction* window to create the new junction.

Delete Junction

Delete the selected junctions. This button is only enabled if one or more items are selected in the list.

10.2. Creating an HPSS fileset

This section provides information on how to create HPSS filesets. Only the HPSS root user and SSM principal are allowed to create filesets.



In order to successfully perform fileset administration, the DB2 Helper Program must be bound. See the Generate and bind the DB2 Helper program section of the HPSS Installation Guide (in the HPSS installation and infrastructure configuration chapter) for more information.

*An HPSS fileset can be created by using the Create Fileset SSM window, or by using the **crtfset** utility. Refer to the **crtfset** man page for more information.*

SSM automatically generates a fileset ID for the new fileset.

After a fileset has successfully been created, a junction will need to be created to get to the new fileset name space from the root node fileset name space.

10.2.1. Create Fileset window

The *Create Fileset* window is used to create new HPSS filesets.

To open the window, from the *HPSS Health and Status* window **Monitor** menu select **Filesets & Junctions** to display the *Filesets & Junctions List*, then click the **Create Fileset** button.

After creating a new HPSS fileset, a new junction will be needed to connect this new fileset to the HPSS name space. See *Section 10.5, “Creating a junction”* for more information.

Field descriptions.

Fileset Name

The name to be assigned to the fileset. This name must be unique to the realm in which HPSS resides.

Fileset State

The state of the fileset. If **Read** is ON, the fileset will be available for reading. If **Write** is ON, the fileset will be available for writing.

File Family

The name of the file family assigned to this fileset. If the file family is to be other than the default, the file family must have been previously created.

Class of Service

The name of the Class of Service assigned to this fileset. This menu will display the available Classes of Service.



Advice

The **File Family** and **Class of Service** fields are optional. They provide the capability to force all data stored in the fileset to be assigned to a particular file family or Class of Service or both.

User Data

Any name or other data (up to 128 bytes) that the administrator wishes to associate with the fileset. The information can be ASCII or binary, although ASCII is easier to work with in this field. HPSS does not use this field in any way; it is strictly for user convenience in annotating the fileset. The field may be left blank.

The field is displayed as printable ASCII characters where possible. Nonprintable bytes are displayed in backslash-octal notation, where each byte is shown as a backslash ("\") followed by a three-digit octal number. For example, a tab character (decimal value 9) would be displayed as "\011". Backslash characters are themselves displayed as two backslashes ("\\"). Trailing null (that is, zero) bytes are not displayed.

To modify this field, enter data in the same format. Printable characters (except for the backslash) can be entered normally. Backslash characters must be entered as "\\". Nonprintable characters must be entered in backslash-octal notation. You need not specify leading zeros on the octal numbers, EXCEPT when the nonprintable byte is followed by a printable octal digit character (0 - 7). In that case you must specify all three digits of the octal number to distinguish it from the printable character following.

Core Server

The name of the Core Server which will create and manage the fileset.

UID

The user ID identifying the user owning the root node of the fileset.

GID

The group ID identifying the principal group owning the root node of the fileset.

Permissions

The initial UNIX-style permissions to be assigned to the root node of the fileset. There are nine check boxes arranged in a matrix with the columns specifying "r" (read), "w" (write) and "x" (execute) permissions, and the rows specifying the three classes of users to which the permissions apply (User, Group, and Other). If a check box contains a check, it means that access is permitted. For example, if the check box in the "x" column and the Group row contains a check, it means that Group users have execute access.

Buttons.**Create**

Creates the fileset. If the create succeeds, a success message will appear on the status line at the bottom of the window; otherwise an error message will appear.

Dismiss

Closes the window.

10.3. Managing existing filesets

This section describes how to look up information on, modify, or delete filesets.

10.3.1. Core Server Fileset Information window

Core Server Fileset Information

File Edit Help

Fileset ID: 3272683715,,2360846864

Fileset Name: FilesetRoot.3202

Subsystem Name: Subsystem #1

File Family: Not in a family

Class of Service: NONE

User Data: RootOfRoots Fileset

Files: 49

Directories: 118

Sym Links: 0

Hard Links: 0

Junctions: 3

Core Server: Core Server

Fileset State

☒ Read

☒ Write

☐ Destroyed

Permissions

	r	w	x
User	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Group	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Other	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>

UID: 0

GID: 0

Delete ☐ Freeze Update Refresh Dismiss

Retrieving data - Succeeded

This window allows an administrator to view, update, and delete the Core Server information associated with a fileset. This information is acquired from the Core Server. While this window remains open, it will automatically update whenever change notifications about any field (except UID, GID, and permissions) are received from the Core Server.

To open the window, from the *HPSS Health and Status* window **Monitor** menu select **Filesets & Junctions** to display the *Filesets & Junctions List*, select one or more filesets, then click the **Detail** button.

Any changes made to fields on this window are sent directly to the Core Server when the administrator clicks on the **Update** button; the changes are effective upon successful update. To mitigate automatic updates overwriting changes in progress, it would be wise to freeze this window by checking the **Freeze** check box before making changes.

The **Fileset State** can be used to prevent any changes to fileset data or to fileset metadata. Changing the state to just Read (unchecking the **Write** check box) will allow reading, but no changes will be allowed to the data or metadata. Changing the state to Destroyed will prevent both reading and writing.

Field descriptions.

Fileset ID

The ID number which identifies the fileset. A fileset ID is displayed as two double-comma-separated unsigned integer numbers. A new Fileset ID can be entered as two double-comma-separated unsigned integer numbers, as two single-period-separated unsigned integer numbers, as a 64-bit hexadecimal number that begins with *0x*, or as an unsigned 64-bit number.

Fileset Name

The unique name which has been assigned to the fileset. The administrator can change the **Fileset Name** as long as the new name is unique to the realm in which HPSS resides.

Subsystem Name

The subsystem name to which the fileset has been assigned. This is the subsystem of the Core Server which controls the fileset.

File Family

The name of the file family to which the fileset is assigned. If this field contains `Not in a family`, the fileset has not been assigned to a family.

Class of Service

The name of the Class of Service to which the fileset is assigned. If this field contains `NONE`, the fileset has not been assigned to a Class of Service.

Fileset State

This set of check boxes displays fileset states, and allows them to be changed.

Read

If checked, the fileset is available for reading.

Write

If checked, the fileset is available for writing.

Destroyed

If checked, the fileset cannot be written to or updated; in fact, it is not possible to set any of the attributes except the **Fileset State** attribute.

User Data

This field is available for storing up to 128 bytes of data. The information can be ASCII or binary, although ASCII is easier to work with in this window.

The field is displayed as printable ASCII characters where possible. Nonprintable bytes are displayed in backslash-octal notation, where each byte is shown as a backslash ("**") followed by a three-digit octal number. For example, a tab character (decimal value 9) would be displayed as "*\011*". Backslash characters are themselves displayed as two backslashes ("**"). Trailing null (that is, zero) bytes are not displayed.

To modify this field, enter data in the same format. Printable characters (except for the backslash) can be entered normally. Backslash characters must be entered as "**". Nonprintable characters must be entered in backslash-octal notation. It is not necessary to specify leading zeros on the octal numbers, EXCEPT when the nonprintable byte is followed by a printable octal digit character (0 - 7). In that case, all three digits of the octal number must be specified.

Files

The number of files in this fileset.

Directories

The number of directories in this fileset.

Sym Links

The number of symbolic links in this fileset.

Hard Links

The number of hard links in this fileset.

Junctions

The number of junctions in this fileset.

Core Server

The name of the Core Server that handles this fileset.

UID

The user ID identifying the user owning the root node of the fileset.

GID

The group ID identifying the principal group owning the root node of the fileset.

Permissions

The UNIX-style permissions assigned to the root node of the fileset. There are nine check boxes arranged in a matrix with the columns specifying "r" (read), "w" (write) and "x" (execute) permissions, and the rows specifying the three classes of users to which the permissions apply (User, Group, and Other). If a check box contains a check, it means that access is permitted. For example, if the check box in the "x" column and the Group row contains a check, it means that Group users have execute access.

Buttons.**Delete**

Requests deletion of the currently displayed fileset. A confirmation will be requested before the request is sent to the Core Server. If the deletion is successful, a message will appear on the status line at the bottom of the window and all fields and buttons will be disabled (except for the **Dismiss** button). Note that only empty filesets may be deleted.

Freeze

Clicking on this check box freezes the window; in other words, all automatic updates to the window are suspended. Toggling the check box again returns the window, and the check box, to their normal conditions. Any accumulated changes that occurred while the window was frozen are immediately displayed.

Update

Requests that the currently displayed fileset information replace the fileset information currently kept by the Core Server.

Refresh

Requests that the current information about this fileset be fetched from the Core Server.


Dismiss

Closes the window.

10.4. Deleting filesets

Filesets must be empty in order to be deleted. The Core Server will reject the deletion request if the fileset is not empty.

Filesets may be deleted by selecting the filesets on the *Filesets & Junctions List* and clicking the **Delete** button. The following confirmation will be displayed:



To continue with the fileset deletion, click **Yes**.

Filesets may also be deleted from the *Core Server Fileset Information* window using the **Delete** button.


10.5. Creating a junction

Only the HPSS root user and SSM principal are allowed to create junctions.

A junction is a name space object that points to a fileset and is similar to a persistent UNIX mount point. The fileset pointed to may reside in another subsystem.

Junctions can be created using SSM or by using the utility routine **crtjunction**. For more information, refer to the **crtjunction** man page.

10.5.1. Create Junction window



This window allows you to create a junction named by the **Absolute Junction Path Name** to the specified fileset. HPSS defines a junction as a name space object that *points* to a directory. A fileset is

an independent directory subtree, managed as a unit. The fileset to which the junction will point may be managed by the same Core Server where the junction is being created or it may be managed by another HPSS Core Server. HPSS junctions are similar to persistent UNIX mount points.

Although HPSS allows junctions to any directory, this window can only be used to create junctions to directories that are also filesets.

Field descriptions.

Fileset ID

The ID number which identifies the fileset to which the junction will point. This field is filled in from the Fileset ID selected on the *Filesets & Junctions List*.

Fileset Name

The unique name which identifies the fileset to which the junction will point. This field is filled in from the Fileset Name selected on the *Filesets & Junctions List*.

Absolute Junction Path Name

The absolute path name, in the HPSS name space, of the junction that is to be created. If the path name to the junction is not valid, the creation will fail.

Buttons.

Create

Creates the junction. If the create succeeds, a success message will appear on the status line at the bottom of the window; otherwise an error message will be displayed. At this point, data can be entered for another junction, or the window can be dismissed.


Dismiss

Closes the window.

10.6. Deleting a junction

Junctions can be deleted using SSM or by using the utility routine **deljunction**. For more information, refer to the **deljunction** man page.

To delete a junction using SSM, select the junctions to be deleted from the *Filesets & Junctions List* and click the **Delete Junction** button. You will be asked to confirm the deletion with the following dialog:



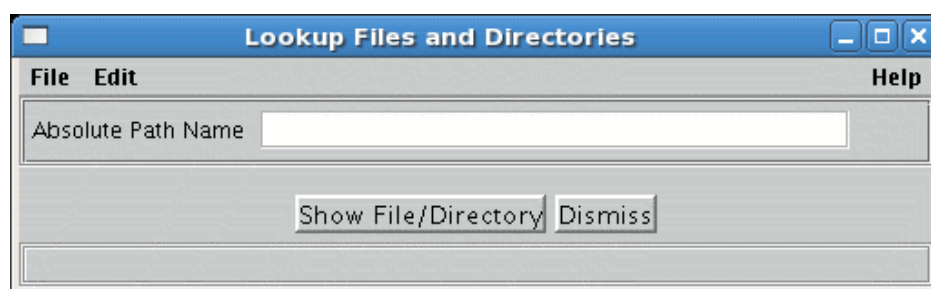
To continue with the junction deletion, click **Yes**.

Chapter 11. Files, directories and objects by SOID

This chapter describes two ways to display basic information about files and directories stored in HPSS. Starting with a fully qualified path name, you can look up a file or directory in the system and display information about it. In other cases, your starting point may be a SOID, a Storage Object ID, which is the internal computer generated name for files, directories, and virtual volumes. Some utility programs report the names of objects of interest as SOIDs.

SOIDs are more difficult to deal with manually as they are not intended to be used that way. But they contain all the necessary information to identify the server, object, and object type, so they can be used to zero-in on a very specific internal data structure, such as a bitfile or virtual volume, that would otherwise be very difficult to isolate and display.

11.1. Files and Directories window



This window is reached through **Monitor > Lookup HPSS Objects > Files & Directories**. The pathname of either an HPSS file or directory may be entered. Use the **Show File/Directory** button to display specific information.

Field descriptions.

Absolute Path Name

The absolute path name (or fully qualified path name), in the HPSS name space, of the file or directory that is to be displayed. If the path name to the file or directory is not valid, then lookup will fail. If the path name is to a symbolic link or junction, then the object displayed will be the object pointed to by the symbolic link or junction.

Buttons.

Show File/Directory

If the Absolute Path Name exists, clicking on this button will open the *File/Directory Information* window. The *Lookup Files and Directories* window will remain open so that the administrator may specify another file or directory.

Dismiss

Closes the window.

11.1.1. File/Directory Information window

The screenshot shows a window titled "File/Directory Information (on hpss-dev02.ccs.ornl.gov)". The window has a menu bar with "File", "Edit", and "Help". The "Path Name" field contains ".Trash/file01.1061249".

The main content area is divided into several sections:

- Object Information:**
 - Object Type: File
 - Class of Service: 1 | Class of Service 1
 - File Family: Not in a family
 - Subsystem Name: Subsystem #1
 - Realm Id: 10000
 - Account: 0
 - Read Count: 0
 - Write Count: 0
 - Link Count: 1
 - Creation Time: Nov 6, 2014 8:56:45 AM
 - Modify Time: Nov 6, 2014 8:58:49 AM
 - Last Written Time: (empty)
 - Last Read Time: Nov 6, 2014 8:59:04 AM
 - Data Length: 0
 - UID: 0
 - GID: 0
 - Permissions: rw-rw-r--
- Option Flags:**
 - ☐ Don't Purge
- Extended ACL:**
 - Object ACL ☐ IC ACL ☐ IO ACL ☐
 - Extended ACL ☐
- Fileset Information:**
 - Fileset Id: 3272831213,,3838938124
 - Fileset Root Object Id: 1
- Fileset State:**
 - ☒ Read ☒ Write ☐ Destroyed
- Trashcan Information:**
 - Parent Id: 3
 - User Id: 0
 - Time Deleted: Nov 6, 2014 8:58:49 AM
 - Time Created: Nov 6, 2014 8:56:45 AM
 - Time Last Read: Nov 6, 2014 8:56:45 AM
 - Time Modified: Nov 6, 2014 8:58:49 AM
 - Original Pathname: ./home
 - Original Filename: file01
- Bitfile Id:**
 - Object Id: 00000001-07-0000000d-01e465bcbe69dbec-0ba2
 - Hash Value: 3200 ☒ Valid
- Handle to Object:**
 - Object Id: 1,061,249
 - Object Hash: 9829
 - File Id: 1,061,249
 - File Hash: 9829
 - Type: File
 - Flags: 0x0
 - Generation: 8,215,264,866,564,785,059
 - Core Server Id: 13

At the bottom of the window, there are three buttons: ☐ Freeze, Refresh, and Dismiss.

This window shows details about a file or directory.

Field descriptions.

Path Name

The pathname to either the file or the directory.

Object Type

The type of object being displayed, either File or Directory.

Class of Service

The name of the Class of Service in which the file is stored. If the displayed object is a directory the value of this field will be `NONE`.

File Family

The name of the file family to which the file has been assigned. If the file has not been assigned to a family, the value of this field will be `Not in a family`.

Subsystem Name

The name of the HPSS subsystem which contains the file or directory.

Realm Id

The ID number which identifies the realm which encompasses the file or directory.

Account

The account index that is associated with the file or directory. This value is used by the HPSS accounting subsystem.

Read Count

This field applies only to files and is the number of read operations that have been issued against the file. Reading a file through most HPSS interfaces may cause multiple read operations to be issued, so this value may increase by more than one with each user read of a file.

Write Count

This field applies only to files and is the number of write operations that have been issued against the file. Writing a file through most HPSS interfaces may cause multiple write operations to be issued, so this value may increase by more than one with each user write of a file.

Link Count

For directories, this is the count of the number of directories that point to this directory. This includes "." (current directory) and ".." (parent directory).

Creation Time

Time and date that the file or directory was created.

Modify Time

Time and date that the metadata associated with the file or directory was last modified.

Last Written Time

This field only applies to files and is the time and date of the last write to the file. If this field is blank, the file has never been written.

Last Read Time

Time and date of the last read from the file, or the time and date of the last read from the directory.

Data Length

For directories, this is the byte length of the metadata associated with this directory. For files, this is the largest numbered byte in the file. Note that this is not necessarily the number of bytes in the

file as files can be written sparsely with gaps. Also, it is not necessarily the highest written byte, since the Core Server supports POSIX `clear` and `truncate` commands which can leave a gap at the end of the file.

UID

The user ID of the owner of the file or directory.

GID

The group ID of the owner of the file or directory.

Permissions

Permissions granted to the file or directory. These are displayed in the standard UNIX mode format (read, write, execute permissions for user, group, and other). These are base permissions and do not take into account the effects of any ACLs on the file or directory.

Option Flags

Don't Purge

This field applies only to files. If the **Don't Purge** check box is selected, the file is purge locked and will not be purged from the top level of the hierarchy. If the **Don't Purge** check box is not selected, the file may be purged as usual.

Super Don't Purge

This field applies only to files. **Super Don't Purge** represents an system, non-user purge lock. It is usually associated with a recover operation. If **Super Don't Purge** is checked, it cannot be cleared except by cancelling the system operation which has set the flag (e.g. recover).

Purge On Migrate

This field applies only to files. When this field is selected, any time the file is migrated, it is immediately purged from a source disk level. This field does not take effect if the file is purge locked.

Extended ACL

If any ACL entry other than the default ACL entries exist, then the file or directory is said to contain extended ACLs. There are three type of ACLs that could have extended ACLs:

Object ACL

HPSS Name Space Object ACL

IC ACL

HPSS Name Space Initial Container ACL

IO ACL

HPSS Name Space Initial Object ACL

A check mark will be put into each of the ACLs containing extended ACL entries for this file or directory. For example, if a directory contains an USER ACL entry in the Object ACL and Initial Container ACL, then both the **Object ACL** and **IC ACL** check boxes will contain check marks.

Fileset Information

The Fileset Information group contains key fileset attributes for the root node of the fileset containing the file or directory.

Fileset Id

The ID number which identifies the fileset containing the file or directory . The 64-bit Fileset ID is displayed as two 32-bit unsigned integers separated by two commas.

Fileset Root Object Id

The ID number which identifies the root node of the fileset which contains the file or directory. This value is a 64-bit object ID.

Fileset State

Filesets have three access states: read, write, and destroyed. If a fileset allows reading, a check mark is displayed in the box labeled Read. If a fileset allows data to be written, a check mark is displayed in the box labeled **Write**. If a fileset is turned off, a check mark is displayed in the box labeled **Destroyed**.

Trashcan Information

When the HPSS trashcan feature is enabled, each HPSS object that is deleted causes a record (row) to be made in the NSTRASH database table. If the object that is displayed is currently in the trashcan, then the **Trashcan Information** fields will be filled in with the state of the object in the trash. If the object being displayed is not currently in the trashcan, then the **Trashcan Information** fields will be empty.

Parent Id

The Object ID of the Parent directory that contained the deleted entry.

User Id

User ID of the person who actually performed the deletion. (It isn't always the owner who performs the deletion.)

Time Deleted

The epoch time that the deletion was performed. This time is needed so that we can compute the length of time that an entry has been in a trashcan.

Time Created

At the time an entry is deleted this time is taken directly from the entry's object record. If this entry is ever undeleted, this time will be restored back into the entry's object record.

Time Last Read

At the time an entry is deleted this time is taken directly from the entry's object record. If this entry is ever undeleted, this time will be restored back into the entry's object record.

Time Modified

At the time an entry is deleted this time is taken directly from the entry's object record. If this entry is ever undeleted, this time will be restored back into the entry's object record.

Original Pathname

The full pathname to the deleted entry. Note that Path is always relative to the root node of the fileset which contained the deleted entry. Path is useful when attempting to undelete an entry—it allows us to try to move the file back to its original location. It can also be useful when simply attempting to recall the original location of a particular trashcan entry.

Original Filename

This is the original entry name before the entry was moved into the trashcan. When entries are moved into trashcans their Names are purposely changed. This is done to prevent possible

name conflicts. The Names are changed using the following format: <Name>.<ObjectId>. For example, if file `Foo` is deleted, its Name in the trashcan might be `Foo.12345`. Note that if an attempt is made to undelete `Foo.12345`, its name can be restored to its original Name, `Foo`.

Bitfile Id

If the object is a File, the bitfile ID to that file is displayed here. This **Bitfile Id** is used by the Trashcan Statistics Thread to obtain and report the length of all Files residing in trashcans. If the object is not a File, the **Bitfile Id** will not be displayed on the screen.

Object Id

Every bitfile is assigned a unique identifier.

Hash Value

The **Hash Value** is used to identify the database partition to which the bitfile is assigned. If the **Valid** check box is not set, the value of **Hash Value** is not available.

Valid

This check box is set if the value of **Hash Value** is available and not set otherwise.

Handle to Object

Object Id

Every file, directory, symlink, and other name space object is assigned a unique identifier. The **Object Id** field in this object handle contains this unique identifier.

Object Hash

The **Object Hash** is generated by running the object's **Parent Id** and the object's Name through a hashing algorithm. This hash value is then used to identify which database partition the object is assigned to.

File Id

File objects need an additional identifier when the file is a HardLink. If the File object is not a HardLink, the **File Id** is equal to the **Object Id**.

File Hash

The **File Id** and the **File Hash** fields are only used when the object identified by the **Object Id** field is a HardLink. When this is the case, the **File Hash** identifies the database partition that holds the metadata information needed by the HardLink.

Type

The **Type** of the object identified by this object handle. Valid object Types are File, Directory, Junction, SymLink, and HardLink.

Flags

If this object is the root node of a fileset, the `NS_OH_FLAG_FILESET_ROOT` bit is set in the **Flags** field.

Generation

The **Generation** field contains a uniqueifier. When verifying this object handle, this uniqueifier is compared to the **Generation** number kept in the object's metadata.

Core Server Id

This is the ID of the Core Server that issued this object handle.

Active Writes

Number of active writes against an open file. If the file is not open, the value is -1.

Active Reads

Number of active reads against an open file. If the file is not open, the value is -1.

Active Migrates

Number of active migration operations against an open file. If the file is not open, the value is -1.

Active Modify

Number of active modification operations against an open file. If the file is not open, the value is -1.

Active Copy

Number of active copy operations against an open file. If the file is not open, the value is -1.

Active Stages

Number of active stage operations against an open file. If the file is not open, the value is -1.

Freeze

Clicking on this check box freezes the window; in other words, all automatic updates to the window are suspended. Toggling the check box again returns the window, and the check box, to their normal conditions. Any accumulated changes that occurred while the window was frozen are immediately displayed.

Buttons.**Refresh**

Requests that the current information about this fileset be fetched from the Core Server.

Dismiss

Closes the window.

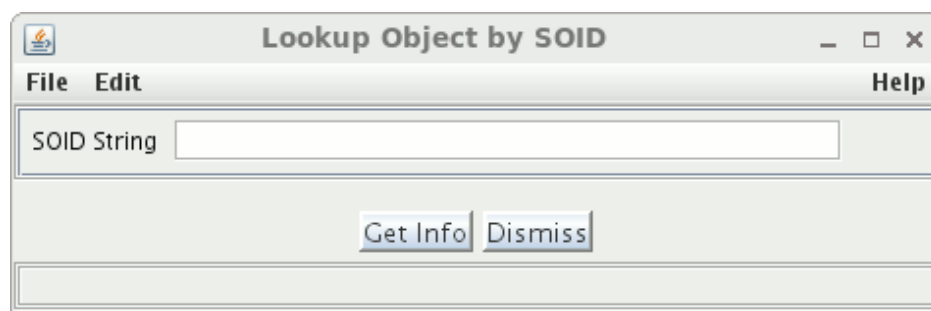
Note: The table below indicates which fields and attributes apply to both files and directories or to just one of the two. Refer to it to see which fields on the *File/Directory Information* window apply to files, directories, or both.

Table 11.1. Attributes of files and directories

Attribute	File	Dir
Account	x	x
BitfileId	x	
RealmId	x	x
Comment	x	x
CompositePerms	x	x
COSId	x	
DataLength	x	x
EntryCount		x

Attribute	File	Dir
ExtendedACLs	x	x
FamilyId	x	
FilesetHandle		
FilesetId	x	x
FilesetRootId	x	x
FilesetStateFlags	x	x
FilesetType	x	x
GID	x	x
GroupPerms	x	x
LinkCount		x
ModePerms	x	x
Name	x	x
OpenCount	x	
OptionFlags	x	
OtherPerms	x	x
ReadCount	x	
SubSystemId	x	x
TimeCreated	x	x
TimeLastRead	x	x
TimeLastWritten	x	x
TimeModified	x	x
Type	x	x
UID	x	x
UserPerms	x	x
WriteCount	x	

11.2. Objects by SOID window



To display this window, select **Monitor** from the *HPSS Health and Status* window, then select **Lookup HPSS Objects**, and then **Objects by SOID**. This window allows you to open an information

window for an HPSS object which you specify by the object's HPSS Storage Object ID (SOID). The object types supported by this screen are Bitfiles and Virtual Volumes (Disk and Tape).

Field descriptions.

SOID String

Storage Object ID string is the internal computer generated name for files, directories, and virtual volumes. HPSS utility programs that display SOIDs should present this field in the correct format.

Buttons.

Get Info

Once all of the fields have been filled in, clicking on this button will open the information window for the object type you selected. The *Lookup Object by SOID* window will remain open so that another SOID may be specified.

Dismiss

Closes the window.

Chapter 12. Tape aggregation

This chapter discusses the following operations:

- Overview of tape aggregation
- Tape aggregation performance considerations
- Ordered Migration and Full Aggregate Recall
- Configuring tape aggregation
- Caveats

12.1. Overview of tape aggregation

Tape aggregation reduces file processing time when migrating relatively small files from disk to tape. In certain situations it may improve very-small-file disk-to-tape migration performance by two orders of magnitude over traditional, file-by-file migration. Tape aggregation also allows more files to be stored on a single tape volume by dramatically reducing the number of tape marks written to tape.

HPSS defaults to writing tape aggregates for all disk-to-tape migrations. If for some reason any file cannot be migrated to a tape aggregate, HPSS may instead migrate it as a single file. The maximum number of files in an aggregate is set to 50,000 files. Ensure the Core Server configuration is set correctly to accommodate the appropriate Core Server I/O limits. See *Section 5.2.2.2, “Core Server I/O limit configuration”* for more detail. In addition, the number of maximum open files configured for HPSS may need to be greater since the maximum number of files allowed per aggregate has increased. See *Section 5.2.2.3, “Core Server max open files configuration”* for more information.

During single file migration, HPSS migrates each file into a separate tape segment. This generally causes the tape drive to flush the file data to the tape and subsequently write a tape mark. It is not unusual for the majority of the time spent migrating small files to be dominated by writing the tape mark.

Tape aggregation is enabled automatically in the disk migration policy, and if there are a sufficient number of migration candidate files ready to be migrated, HPSS will pack the data of many small files into a single tape storage segment. Thus, the 'end of file' tape mark is only written at the end of the batch of files, rather than one tape mark for each file being migrated.

Note that tape aggregation is managed internally by HPSS. It has no effect on file security nor on how files are accessed and generally has little effect on file retrieval performance.

12.2. Tape aggregation performance considerations

In order to get good performance with tape aggregation, there are several points to consider.

First, there needs to be a sufficient number of files ready to migrate. If migration runs too often, or if the HPSS system doesn't tend to ingest a large number of small files in a relatively short period of time, then there may be little or no performance benefit to tape aggregation since there will be very few files, if any, in each aggregate. Overall numbers of tape marks written to tape will still be reduced, but not as dramatically as in situations where there are relatively many files to migrate.

Second, tape aggregation performs best with relatively small files since most of the time spent to migrate them occurs during tape mark processing when the data is flushed to the media. From a data transfer rate performance perspective, it usually isn't useful to aggregate large files since the tape mark processing time is dominated by the time needed to write the file data. However, even with large files, the benefits of reduced tape mark usage still leads to storing more files on a tape volume.

For very large aggregate sizes, HPSS disk Movers may begin timing out during the migration. The Mover environment variable **MVR_CLIENT_TIMEOUT** may need to be increased on the Movers. See *Section 5.2.6.2, "Additional Mover configuration"* for more detail.

The size of an aggregate is configurable by setting the "Max Aggregate Size" field in the disk migration policy. This value is an absolute value, and care must be taken to ensure a correct value if the migration policy is shared between storage classes.



Total Migration Stream Distinction

HPSS 9.2+ supports a larger number of files in an aggregate. This may result in a larger number of bytes being written in a single aggregate. One migration stream consumes $2 \times$ "Max Aggregate Size" bytes before a secondary migration stream becomes active. As a result, 9.2+ systems may use fewer migration streams than earlier versions of HPSS.

12.3. Ordered Migration and Full Aggregate Recall

The Ordered Migration feature orders files being migrated from disk to tape by directory or by create time. Further, for files being aggregated on tape, the tape aggregates can be used to group files according to their directories. This is especially pertinent when considering how to employ the Full Aggregate Recall feature.

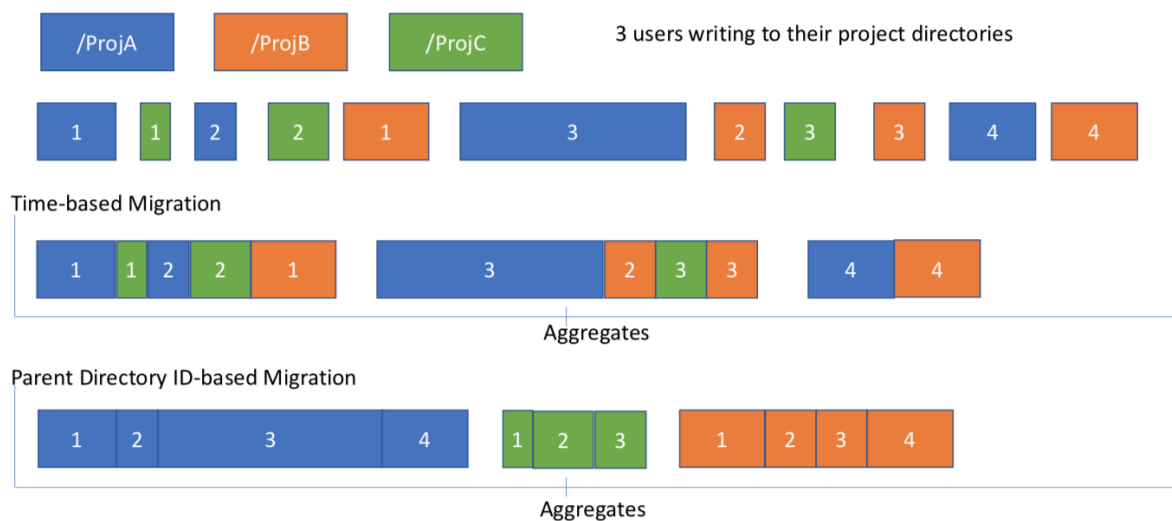
The objective of Ordered Migration is to co-locate data on tape, with either the directory of the files or the files' create times being the co-location criteria. A site may want to combine this feature with Full Aggregate Recall to potentially minimize the number of tape operations required to bring a directory's files back from tape. Of course, the migration policy will continue to honor the migration policy settings for when a file is eligible for migration and how often migration will run.

Ordered Migration also makes use of a **Disk Migration Policy** configuration option which specifies whether or not aggregates should be optimized for reading or writing. **Read-Optimized Aggregates** is defined as having only one ordering key value per aggregate (e.g. only files contained in Directory X). Select **Read-Optimized Aggregates** if you want this behavior. When **Read-Optimized Aggregates** is not selected, HPSS will create write-optimized aggregates: aggregates that may contain files from more than one ordering key value (e.g. files contained in directory X followed by files contained in directory Y, and so on). This is the default behavior. It is another important aspect to consider when preparing to employ Full Aggregate Recall.

Full Aggregate Recall allows an entire aggregate to be streamed from tape to disk. When this feature is on, and a stage is requested for a file which resides within a tape aggregate, then (apart from deleted files and already-staged data) the entire aggregate is staged to disk via a minimal number of tape I/O operations.

Migration Order and **Full Aggregate Recall** are configured on a per Class of Service basis. Refer to *Section 6.3.2, “Class of Service Configuration window”* for details on configuring the Class of Service.

12.3.1. Ordered migration example - read-optimized aggregates



12.4. Configuring tape aggregation

Tape aggregation is turned on by default. To disable tape aggregation and migrate as single file, set the "Max Aggregate Size" field to 0 in the migration policy.

12.5. Debugging tape aggregation

If files are not being migrated into aggregates in the expected manner, it can be helpful to turn on batch migration debugging. To do this, add the environment variable `HPSS_BATCH_MIGR_DEBUG_LEVEL=1` to the `/var/hpss/etc/env.conf` file. Then restart or reinitialize the Core Server to make it notice the new environment variable.

The Core Server will then issue additional DEBUG level log messages when certain aggregation errors are encountered, such as a file being rejected from a batch. These log messages will include the string "BATCH-FILE-DEBUG."

To stop the extra logging, remove the definition of the environment variable from the `env.conf` file, or reset it to 0. Then restart or reinitialize the Core Server.

12.6. Caveats

Tape-to-tape migration of files in aggregates is an unsupported configuration. If multiple copies of aggregated files are needed:

- Set the "Max Aggregate Size" field to 0 in the parent disk storage class's migration policy, and
- Configure the corresponding hierarchy for migration from the disk storage class to two or more tape storage classes.

Chapter 13. Tape Ordered Recall (TOR)

This chapter discusses the following operations:

- Overview of TOR
- TOR ordering behavior
- TOR performance considerations
- Configuring TOR
- Monitoring TOR
- Caveats

13.1. Overview of TOR

Tape Ordered Recall reduces file processing time when recalling data from tape. In certain situations it may improve tape recall rates by one to two orders of magnitude over unordered access. TOR works with or without the enterprise tape Recommended Access Order feature.

TOR is an HPSS feature that implements scheduling, queueing, and reordering of tape requests. If there are multiple tape read requests for the same virtual volume in HPSS, those requests will be queued up in the HPSS Core Server in a Virtual Volume Queue. That queue will then be ordered before the next request is pulled off to be worked, so that the requests are worked using the queueing criteria specified in the *Core Server Configuration* window.

Tape does not deal well with random access. Due to the nature of the media, seek times are generally measured in seconds. This makes unordered tape I/O particularly poor performing, especially when the random accesses involve relatively small files.

TOR attempts to reduce tape seek times by issuing tape requests after they have been ordered. TOR is always on, but has a number of configuration options which can impact its behavior and HPSS tape recall performance.

There are a number of APIs which implement batch staging. These APIs are described in detail in the *HPSS Programmer's Reference*. They include support for issuing multiple synchronous or asynchronous stage requests with a single API call and retrieving schedule status along with stage status data. Additionally, the **quaid** tool can take advantage of these APIs to facilitate mass recall from tape. See **man quaid** for details and usage.

Additionally, the **lori** tool makes use of ordered recall APIs similar to staging. The **lori** tool takes advantage of the read queueing APIs. See **man lori** for details and usage.

13.2. TOR ordering behavior

There are two queuing methods available with TOR: **FIFO** and **Offset**. These are described in detail below.

FIFO

FIFO will issue incoming requests in the order they were received by the Core Server tape subsystem. This generally reflects the order that clients submitted the requests.

Offset

Offset orders the requests using offset information on the volume. The mechanism for this can differ depending upon the tape drive that the volume is mounted in. For certain enterprise drives, they will make use of a drive feature called Recommended Access Order (RAO). This allows the tape drive to inform the application of the best order to read a sequence of requests from that particular tape. If the volume is mounted in a drive without RAO support, the Core Server will order the requests in linear tape order, beginning with the current position of the tape.

Other aspects of TOR behavior

1. Ordering occurs when the tape is idle, after a request has come into the scheduler and has not been ordered.
2. Ordering requires that the tape be mounted in order to support the RAO feature.
3. Ordering does not induce artificial wait times for queued items; if an item enters the queue as the sole item, it will immediately run. This means that, upon generating a set of items for recall, the first request to hit the queue will run without being ordered in with the other requests.

13.3. TOR performance considerations

In order to get good performance with Tape Ordered Recall, there are several points to consider.

First, there need to be a sufficient number of parallel stage or read requests coming in from the application for data on tape. If tape read requests come in a handful at a time there may be little to no gain seen with TOR. TOR can only schedule what it knows about, so it is important to provide a steady stream of parallel read requests into HPSS so that TOR may schedule them. Since TOR works at the virtual volume level, it is most beneficial if these parallel reads are for data on a small set of volumes. Direct read from tape, synchronous stages, and asynchronous stages all go through the TOR queuing logic, so the type of read request doesn't impact whether requests will be ordered. More read requests queued up allows for better scheduling, which should result in less seek time and improved tape recall performance.

Second, HPSS has a number of configuration options which can impact the performance and behavior of TOR. The major configuration parameters which can affect TOR performance are described below.

Maximum Connections

The Location Server, Gatekeeper (if configured), and Core Server maximum connections will limit the number of client connections, which could limit incoming tape read requests. This parameter should be set to a value which will allow the desired throughput of tape read requests to the server.

Thread Pool Size

Same as above.

Max Active I/O

The Core Server I/O scheduler has a limit to the number of I/O requests that are processed, subsystem-wide. Each read request that triggers a stage counts as a single I/O. In order for a

sufficient number of parallel tape read requests to get passed down to the tape scheduler, it may be necessary to increase this value. For example, in order for the system to process 100 tape read requests in parallel, this value would need to be at least "100".

Max Active Copy

The Core Server I/O scheduler has a limit to the number of data copy requests that are processed, subsystem-wide. Each stage request, including each individual file within a batch stage request, counts as a single copy. Like *Max Active I/O*, this value needs to be configured appropriately for the expected stage workload. For example, in order for the system to process 100 stage requests, this value would need to be at least "100".

Tape Queue Method

The Core Server's Queue Method has a major impact upon performance with TOR. The default is *Offset*; the Core Server will use RAO for tapes volumes mounted in suitable drives and linear tape ordering for all others. Using *FIFO* should only be done if the site is already ordering all requests from tape.

Tape Queue Active Set Size

The Core Server's Active Set Size controls the number of stage requests, per volume, which are scheduled together. For example, if a volume has 10 pending stage requests, and the *Active Set Size* is set to "5", then the first five requests to enter the queue will be ordered and run. Then the next five requests will be ordered and run. Lower values can impact performance because it limits the requests that the scheduler can consider for generating a well-performing I/O schedule. Higher values can increase the average wait time of individual requests into the scheduler.

Level 0 Disk Storage Class Max Segment Size

For stage requests in a disk over tape hierarchy, asynchronous stage requests will be done in chunks based on the disk storage class *Max Segment Size*. Increasing the disk *Max Segment Size* will result in larger read requests from tape, which will result in better performance. Synchronous stage requests are unaffected. For example, to asynchronously stage a 1000 MB file up to a disk storage class with a 100 MB max segment size, it would require 10 reads from tape.

Gatekeeper Settings

For sites which make use of a Gatekeeper and a site policy which restricts opens or stages, there may be adjustments to make to allow additional requests to come through as part of a large set of stages against a single VV. These adjustments may differ depending upon the implementation of your site policy.

13.4. Configuring TOR

To configure Tape Ordered Recall, open the *Section 5.2.2, "Core Server-specific configuration"* screen. Update the fields as necessary. The fields which directly affect TOR are *Tape Queue Method*, *Tape Queue Active Set Size*, and *Tape Max Reader Wait (seconds)*.

RAO-enabled tape drives may have their RAO capability enabled or disabled in SSM using the tape drive configuration screen. See (*Section 7.1, "Configure a new device and drive"*). This provides the finest granularity capability for enabling or disabling RAO per tape drive type. Drives which have RAO disabled will use offset-based ordering unless FIFO is selected for the volume.

There are three additional methods to modify TOR behavior, even while the system is running, without a Core Server reinitialization. These are meant to be used in diagnostic or temporary error situations and are described below.

Globally Disable Scheduling

By touching a file named `$HPSS_TMP/noschedule` (for example, `/var/hpss/tmp/noschedule`), any Core Servers which can read that file will disable all reordering of requests. Requests will still be queued, but they will not be reordered until the next opportunity after the file is removed.

Globally Disable Recommended Access Order

By touching a file named `$HPSS_TMP/nodrive` (for example, `/var/hpss/tmp/nodrive`), any Core Servers which can read that file will not use the enterprise drive Recommended Access Order (RAO) feature, but will instead do linear offset ordering. Upon removing the file, ordering with RAO will be done the next time the queue would be reordered.

Mark an individual volume as FIFO

Individual volumes can be configured to be read in a FIFO manner via the *Core Server Tape Volume* information window. This can assist in reading certain types of damaged volumes.

Tape Ordered Recall writes a periodic dump file, one per subsystem. This file is dumped roughly every five seconds and contains the current schedule, and statistics in JSON format. This can provide a site with data and metrics regarding requests, including their order within the queue, their length, and how long they have been waiting.

This file is located in `$HPSS_TMP/tape_schedule_<subsystem-id>.json` (for example, `/var/hpss/tmp/tape_schedule-1.json`).

13.5. Caveats

1. Recommended Access Order currently has a hardware limit of 2730 requests that can be ordered. Sites using enterprise drives should be aware that while they can set a larger Active Set Size, the scheduler will cap the number of requests that will be ordered via an RAO-capable enterprise drive at 2730.

Chapter 14. User accounts and accounting

14.1. Managing HPSS users

After the HPSS system is up and running, the administrator must create the necessary accounts for the HPSS users. For a new HPSS user, a Kerberos, LDAP, or UNIX ID (depending on authentication type configured) must exist before the user can access HPSS. In addition, if the HPSS user needs to use SSM, an SSM ID must also be created before the user can use SSM. The SSM ID should be created only for the HPSS administrators and operators.

The HPSS User Management Utility (**hpssuser**) provided with HPSS can be used by the administrator to add, delete, and list the HPSS user IDs. The utility must run as root to acquire the necessary authority to create new KRB, LDAP, UNIX, and SSM IDs. Refer to the **hpssuser** man page for more information on how to invoke the **hpssuser** utility.

14.1.1. Adding HPSS users

The **hpssuser** utility can be used by the administrator to add a UNIX user ID, a KRB user ID, an LDAP user ID, and an SSM user ID to the HPSS if these IDs do not already exist. The **hpssuser** utility can be invoked to simultaneously add multiple types of user IDs for a user, or to add an individual user type. Refer to the **hpssuser** man page for more information.



*Ensure that the Core Server is up and running before adding the user ID. The **hpssuser** utility will not be able to create the user's home directory if the Core Server is not available.*

14.1.1.1. Add all user ID types

The utility can be used to simultaneously add all relevant user ID types for a user. When invoked with the **-all** option, the **hpssuser** utility will consult the system authentication and authorization configuration and add user types which are consistent with the configuration. For example, if the system is configured to use Kerberos authentication and LDAP authorization, a KRB user ID and an LDAP user ID will be added in addition to the UNIX user ID, and SSM user ID types.

Invoke the **hpssuser** utility as follows to add the required user ID for an HPSS user:

```
hpssuser -add <user> -all
```

When invoked, the **hpssuser** utility will prompt the user for any required data. The following is an example of adding all user ID types on a system configured to use Kerberos authentication and LDAP authorization. The **-nohome** option indicates that no user home directory will be created in HPSS.

```
# hpssuser -add user1 -all -nohome
User ID#: 300
Primary group name: hpss
Enter password for user1: *****
```

```
Re-enter password to verify: *****
Full name: Test User
Login shell: /bin/ksh
Unix (local/system) home directory: /home/user1
[ adding unix user ]
[ added unix user ]
[ KADMIN_PRINC unset; using kadmin.local for Kerberos ops ]
[ adding kerberos principal ]
[ added kerberos principal ]
HPSS/LDAP home directory: /
Primary group ID#: 210
[ adding ldap principal ]
[ added ldap principal ]
[ adding ssm user ]
1) admin
2) operator
Choose SSM security level
(type a number or RETURN to cancel):
> 1
[ ssm user added : admin ]
```

14.1.1.2. Add a user ID

The **hpssuser** utility can be used to create HPSS user identities and create keytabs by using the following syntax:

```
hpssuser -add <user> -acct [-keytab <path>]
```

The utility will prompt the user for required data for the configuration. Following is an example of adding a UNIX user ID and creating a UNIX keytab:

```
# hpssuser -add user1 -acct -keytab user1.keytab
User ID#: 300
Primary group name: hpss
Enter password for user1: *****
Re-enter password to verify: *****
Full name: Test User
Login shell: /bin/ksh
Unix (local/system) home directory: /home/user1
[ adding unix user ]
[ added unix user ]
[ adding unix keytab entry to 'user1.keytab' ]
[ added unix keytab entry to 'user1.keytab' ]
```

*Ensure that the Core Server is up and running before adding the user ID. The **hpssuser** utility will not be able to create the user's home directory if the Core Server is not available.*

14.1.1.3. Add an SSM user ID

The **hpssuser** utility creates an SSM ID for a user by adding an entry for him to the AUTHZACL table. Refer to *Section 2.3, “SSM user security”* for more information on SSM user security.

Invoke the **hpssuser** utility as follows to add an SSM user ID:

```
hpssuser -add <user> -ssm
```

The utility will prompt the user for the required data. Following is an example of adding an SSM user ID:

```
# hpssuser -add user1 -ssm
[ adding ssm user ]
1) admin
2) operator
Choose SSM security level
(type a number or RETURN to cancel):
> 1
[ ssm user added : admin ]
```

14.1.2. Deleting HPSS users

The **hpssuser** utility can be used by the administrator to delete existing user IDs for an HPSS user. The utility can be invoked to delete all user IDs for the user or to delete an individual ID.

The utility will prompt the user for the required data. Following is an example of deleting the user IDs for an HPSS user on a system configured to use Kerberos authentication and LDAP authorization:

```
# hpssuser -del user1 -all
[ deleting ssm user ]
[ SSM user deleted ]
[ deleting ldap principal ]
[ deleted ldap principal ]
[ deleting kerberos principal ]
[ KADMIN_PRINC unset; using kadmin.local for Kerberos ops ]
[ deleted kerberos principal ]
[ deleting unix user ]
```

14.1.3. Listing HPSS users

The **hpssuser** utility can be used by the administrator to list all existing HPSS user IDs. The utility can be invoked to list all HPSS user IDs or a particular type of user ID.

Following is an example of listing the *user1* user ID on a system configured to use UNIX authentication and authorization.

```
# hpssuser -list user1 -all
[ Kerberos not configured - skipping... ]
[ LDAP not configured - skipping... ]
[ unix(linux) user info ]
uid=300
gid=210
home=/home/user1
shell=/bin/ksh
fullname=Test User
[ ssm user info ]
(admin -> perms = 'rwxcidt')
(operator -> perms = 'r--c--t')
SSM client security ACL entry:
rwxcidt - user - 300 (user1) - 10000 (HPSS_REALM.NAME)
```



Listing all UNIX users is not supported.

14.1.4. Create an SSM client package

The **hpssuser** utility can be used by the administrator to generate an SSM client package (all files required to start an SSM client on another node). The utility also has the ability to include HPSS documentation in the package if the **-packagedoc** option is specified. The following is an example of creating an SSM client package without the HPSS documentation on a system configured to use Kerberos authentication. For the case of a system configured to use UNIX, then `krb5.conf` won't be included in the package. Note that if `krb5.conf` does not exist on a system that's configured to use krb auth, then the command will fail.

```
# hpssuser -ssmclientpkg /tmp/ssmclientpkg.tar
[ packaging ssm client ]
[ creating /tmp/ssmclientpkg.tar ]
ssm.conf
krb5.conf
hpssgui.pl
hpssgui.vbs
hpss.jar
[ packaged ssm client in /tmp/ssmclientpkg.tar ]
```

14.2. Accounting

HPSS maintains accounting information on the usage of the system whether the site charges for usage or not. Sites are encouraged to use the accounting information, even if they do not charge users, to gain a better understanding of the usage patterns of their system.

The accounting policy defines how the usage of HPSS resources is reported. An accounting policy is required whether the site actually charges users for HPSS usage or not.

The accounting policy may be created using the *Accounting Policy* window (select this window by clicking **Accounting** on the **Policies** submenu of the **Configure** menu on the *HPSS Health and Status* window). After the accounting policy is created, it can be viewed, updated, or deleted through the same window. Note, however, that once the accounting style is chosen, it may not be modified.

14.2.1. Accounting Policy window

The screenshot shows the 'Accounting Policy' window. It has a title bar with 'Accounting Policy' and menu items 'File', 'Edit', and 'Help'. The main area contains the following fields:

- Accounting Style:** A dropdown menu with 'SITE' selected.
- Storage Unit Size:** A dropdown menu with 'Bytes' selected.
- Pathname of Executable (UNIX):** A text field containing '/opt/hpss/bin/hpss_acct'.
- Report File (UNIX):** A text field containing '/var/hpss/acct/acct_report'.
- Ordering:** A dropdown menu with 'Account, User, Group' selected.

Below these fields is an 'Options' section with three checkboxes:

- ☐ Account Validation
- ☐ Require Default Account
- ☐ Account Inheritance

At the bottom of the window are four buttons: 'Update', 'Delete', 'Start Over', and 'Dismiss'.

This window allows an administrator to manage the accounting policy. If the **Update** button is visible, then the accounting policy already exists; otherwise, an **Add** button will be displayed. If changes are made to the **Storage Unit Size**, **Pathname of Executable (UNIX)**, **Report File (UNIX)**, or **Ordering** fields, after clicking **Update** the changes will take effect on the next generation of an accounting report. If any of the **Options** flags are changed, clicking the **Update** button and then clicking **Apply Config** on the *HPSS Health and Status* window will reinitialize the Core and Gatekeeper servers to allow the changes to take effect.

An accounting policy is required whether the site actually charges users for HPSS usage or not.

Field descriptions.

Accounting Style

The style of accounting that is used by the entire HPSS system. Valid values are SITE or UNIX. The default value is UNIX.

Under UNIX-style accounting, resource usage is reported by user ID (UID). Each user is allowed to use only one account ID, which has the same numerical value as his user ID. Under SITE-style accounting, each user may use multiple account IDs.

Note that if you use SITE-style accounting, you must either enable **Account Validation** in the accounting policy or select LDAP as your authorization mechanism. See below for a description

of the **Account Validation** field in the accounting policy. See *Section 2.1, "Security services"* for a description of LDAP and other HPSS authorization mechanisms. See the *Accounting policy and validation* section of the *HPSS Installation Guide* (in the *HPSS planning* chapter) for a complete description of accounting styles and account validation.



Advice

This field must be set up properly before any files are written to the system. Under UNIX-style accounting, HPSS obtains the account number from the user's UNIX UID. SITE-style accounting allows the account ID to be set independently of the UNIX UID. Once the accounting policy is configured, the accounting style cannot be changed.

Storage Unit Size

The integral units to be displayed in the report file. The valid values for this field are Bytes, Kilobytes, Megabytes, and Gigabytes. The default value is Bytes.

Pathname of Executable (UNIX)

The UNIX path name of the accounting utility executable. This field may contain any valid UNIX path name. The default value is `/opt/hpss/bin/hpss_acct`.

Report File (UNIX)

The UNIX pathname where the generated accounting report will be stored. The Report File (UNIX) field can contain any valid UNIX pathname. The default value is `/var/hpss/acct/acct_report`.

Ordering

Accounting may report users based on their Account ID, User ID, and/or Group ID. This field specifies which of the identifiers to use when creating an accounting report as well as the order of priority for sorting. For example, when selecting "User, Group", users will be sorted by User ID first followed by Group ID second. Any users that are identical with respect to the chosen identifiers are combined into a single row. For example, if there are two users (Account ID = 1, User ID = 2, Group ID = 3) and (Account ID = 2, User ID = 2, Group ID = 3). If the ordering is set to "User, Group", the information about these two users are combined as they are the same user with respect to the chosen report ordering.

Options

Account Validation

A flag that indicates whether or not authorization of user accounts should be performed by a Gatekeeper. The default value is OFF.



Advice

If you turn this flag ON, you must configure at least one Gatekeeper which will perform the account validation service.

Require Default Account

A flag that indicates whether or not users must have a valid default account index before they are allowed to perform any operation. The default value is OFF. It is only used if Account Validation has been enabled. If this flag is disabled, validation will occur only during file manipulation operations.

Account Inheritance

A flag that indicates whether or not newly created files and directories should automatically inherit the account index used by their parent directory. The default value is OFF. It is only used if Account Validation has been enabled and SITE-style accounting has been selected. If this flag is disabled, new files and directories have the user's current session account index applied to them.

14.2.2. Accounting reports and status

14.2.2.1. Generating an accounting report

The HPSS accounting report tool accumulates accounting information from metadata and generates a report file. The accounting report will include:

- Total file accesses, total number of files, and total space used per account per Class of Service.
- Total file accesses and amount of data transferred per account per Class of Service per storage class.

The accounting report represents a blurred snapshot of the system storage usage as it existed during the accounting run; that is, the accounting information is gathered while the system runs, so the accounting data is changing underneath it, and it may not represent what would be collected if no users were active during the accounting process. Accounting information is maintained per storage subsystem and the accounting reports summarize information for that storage subsystem.

Before creating an accounting report, you must create and fully configure an accounting policy. The initial accounting policy should be set up before any files are created in the HPSS system. Except for the accounting style, the accounting policy fields can be updated any time after the initial setup.

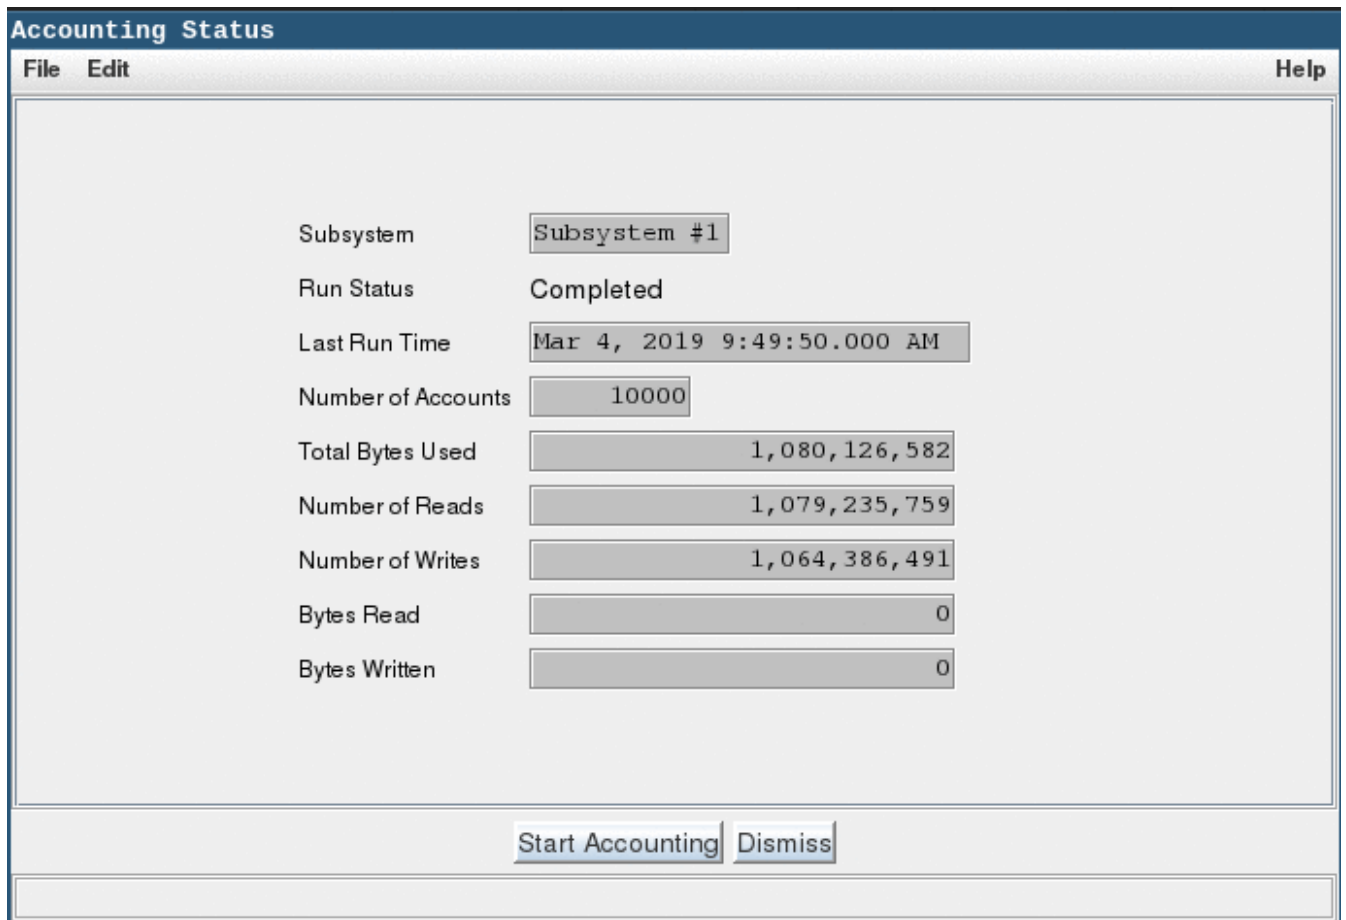
There are two ways to generate an accounting report: from the *Subsystems* window or from the *Accounting Status* window. For each method, begin by selecting the **Accounting Report** menu item from the **Operations** menu on the *HPSS Health and Status* window. This will bring up the *Subsystems* window. Click on the storage subsystem for which an accounting report will be generated. This will highlight the **Accounting Status** and **Start Accounting** buttons.

To create an accounting report, choose either of these methods:

- On the *Subsystems* window, click the **Start Accounting** button. The accounting report should run quickly and a completion status will appear at the bottom of the *Subsystems* window.
- Bring up the *Accounting Status* window by clicking on the **Accounting Status** button on the *Subsystems* window. On the *Accounting Status* window, click the **Start Accounting** button. The accounting report should run quickly and a completion status will appear at the bottom of the *Accounting Status* window.

For either method, the *Accounting Status* window will display the overall statistics from the last accounting run.

14.2.2.2. Accounting Status window



The screenshot shows a window titled "Accounting Status" with a menu bar containing "File", "Edit", and "Help". The main area displays accounting data for "Subsystem #1". The data is organized into two columns: labels on the left and values on the right. The values are displayed in text boxes. At the bottom, there are two buttons: "Start Accounting" and "Dismiss".

Label	Value
Subsystem	Subsystem #1
Run Status	Completed
Last Run Time	Mar 4, 2019 9:49:50.000 AM
Number of Accounts	10000
Total Bytes Used	1,080,126,582
Number of Reads	1,079,235,759
Number of Writes	1,064,386,491
Bytes Read	0
Bytes Written	0

Start Accounting Dismiss

This window allows an administrator to view the accounting status and start accounting.

Field descriptions.

Subsystem

The name of the storage subsystem containing this accounting status data.

Run Status

Current status of accounting run. Possible values are:

- Never run
- Running
- Failed
- Completed
- Report generated

Last Run Time

If accounting is currently running, this is the time the run started; otherwise, it is the time the last run completed.

Number of Accounts

Total number of accounts in the storage subsystem. Set after a successful run.

Total Bytes Used

Total number of bytes accounted for in the storage subsystem. Set after a successful run.

Number of Reads

Total number of reads performed in the storage subsystem during the accounting period covered by the last successful accounting run. Set after a successful run.

Number of Writes

Total number of writes performed in the storage subsystem during the accounting period covered by the last successful accounting run. Set after a successful run.

Bytes Read

Total bytes read from the storage subsystem during the accounting period covered by the last successful accounting run. Set after a successful run.

Bytes Written

Total bytes written to the storage subsystem during the accounting period covered by the last successful accounting run. Set after a successful run.

Buttons.**Start Accounting**

Causes the accounting utility to start running for the storage subsystem. As it runs, the status fields will be updated after the statistics have been changed by the accounting utility.

14.2.2.3. Interpreting the accounting report

The accounting program for HPSS generates a JSON file that contains eight members.

"Date Reported"

The date when the report completed.

"Subsystem"

The subsystem for which the report was run.

"Number of Accounts"

Total number of accounts identified during the accounting run.

"Storage Unit Size"

The size of the units reported in the "Capacity Report", "Bandwidth Report", and "Bandwidth Summary".

"User Tuple"

The ordering used for the report. Each token in this string will be the key to each member in the "Capacity Report", "Bandwidth Report", and "Bandwidth Summary".

"Capacity Report"

An array of objects, each representing accounting information for a unique "User Tuple". The objects will each have the members listed in the "User Tuple" as well as:

"COS"

The class of service.

"Last Mod Time"

The last time this user read/wrote the COS.

"Files Stored"

Total number of files owned by this user in this COS.

"Units Stored"

Total number of units of said "Files Stored".

"Bandwidth Report"

An array of objects, each representing accounting information for a unique "User Tuple". The objects will each have the members listed in the "User Tuple" as well as:

"COS"

The class of service.

"SClass"

The storage class.

"Last Mod Time"

The last time this user read/wrote the COS/SClass.

"Number of Reads"

Number of individual read operations performed on this COS/SClass. In general, reads are counted against the user accessing the file. If a COS is configured to stage bitfiles on open, then all reads will occur in the SClass at the top of the hierarchy.

"Number of Writes"

Number of individual write operations performed on this COS/SClass. In general, writes are counted against the user accessing the file.

"Units Read"

Total number of units read from this COS/SClass.

"Units Written"

Total number of units written from this COS/SClass.

"Bandwidth Summary"

An array of objects, each representing accounting information for a unique "User Tuple". This summarizes the information in "Bandwidth Report" for each "User Tuple" and COS. Thus, this is identical to the "Bandwidth Report" minus the "SClass" field. With the exception of the "Last Mod Time" member, the members are summations of the information from the "Bandwidth Report" for each COS.

Note: File transfer information is not always accounted for when it occurs through interfaces other than the HPSS Client API.

Example accounting report file:

```
{
  "Date Reported": "Fri Mar  8 12:38:04 2019\n",
  "Subsystem": 1,
  "Number of Accounts": 1,
```

```
"Storage Unit Size": 1,
"User Tuple": "AccountId, UID, GID",
"Capacity Report": [
  {
    "AccountId": 2,
    "UID": 3,
    "GID": 4,
    "COS": 5,
    "Last Mod Time": 12345,
    "Files Stored": 2,
    "Units Stored": 800
  }
],
"Bandwidth Report": [
  {
    "AccountId": 2,
    "UID": 3,
    "GID": 4,
    "COS": 5,
    "SClass": 6,
    "Last Mod Time": 7054,
    "Number of Reads": 100,
    "Number of Writes": 53,
    "Units Read": 8024,
    "Units Written": 80000
  },
  {
    "AccountId": 2,
    "UID": 3,
    "GID": 4,
    "COS": 5,
    "SClass": 1,
    "Last Mod Time": 12345,
    "Number of Reads": 42,
    "Number of Writes": 91,
    "Units Read": 103,
    "Units Written": 623
  }
],
"Bandwidth Summary": [
  {
    "AccountId": 2,
    "UID": 3,
    "GID": 4,
    "COS": 5,
    "Last Mod Time": 12345,
    "Number of Reads": 142,
    "Number of Writes": 144,
    "Units Read": 8127,
    "Units Written": 80623
  }
]
}
```

Sites may also wish to correlate the accounting report with a site-provided Account Map (see *Section 14.2.3, “Accounting procedures”*) to determine the appropriate accounting charges.

The HPSS accounting report and a copy of any current site-defined accounting configuration files, if used, should be named with the date and time and stored for future reference. Site administrators

may choose to write scripts to copy or archive the generated accounting report file from its original location.

14.2.3. Accounting procedures

The amount of work needed to support accounting varies according to the style of accounting used, whether or not account validation is being used, and whether the site requires additional information beyond that maintained by HPSS.

For UNIX-style accounting, with or without account validation, no additional work is needed to manage accounting.

For SITE-style accounting, the site must define and maintain a mapping of user IDs to account IDs. For each user, the mapping must define the user's default account ID and any other accounts for which the user is authorized.

- If Account Validation is enabled:
 - The user ID to account ID mapping is maintained in the Account Validation table. This table can store the default account ID of each user and any other accounts for which the user is authorized. The site must maintain this table. HPSS provides the table and a tool for maintaining it, the Account Validation Metadata editor (**hpss_avaedit**). See the **hpss_avaedit** man page for more details.
 - If the site requires additional information beyond that supplied by the Account Validation table, it must create and maintain a site-defined Account Map. See *Section 14.2.3.1, "Site-defined accounting configuration files and procedures"*.
 - For account validation to work, a Gatekeeper must be configured. Also, the **Account Validation** flag must be enabled on the *Accounting Policy* window.
- If Account Validation is disabled:
 - The site will need to create a local site Account Map to maintain a list of the account IDs for which each user is authorized. This is a locally designed and written table, not supported by HPSS. See *Section 14.2.3.1, "Site-defined accounting configuration files and procedures"* for suggestions for designing and maintaining an Account Map.
 - Administrators will need to manage user default accounts by updating the user's LDAP registry information (that is, the "AA=<default-acct-idx>" string in the user's hpssGECOS field).

14.2.3.1. Site-defined accounting configuration files and procedures

This section describes situations in which a site would need additional tables or procedures beyond those provided by HPSS. It offers suggestions on ways to implement these tables and procedures.

Site-defined Account Maps

If a site does not use account validation or if the site requires additional information beyond that maintained in the Account Validation table, it will need to create a local site-defined Account Map.

For example, some sites may need to keep track of a department or company name. The Account Map is designed, created, and maintained by the site, not by HPSS.

The HPSS Account Index of each user should correspond to an Account Map entry that has the site information. The following are examples of possible Account Maps:

SITE-style Account Map:

Acct	User	UID	Charge	Update_flag
12	dlk	3152	5A12x401	0
27	dlk	3152	5A12x501	0
341	dlk	3152	5A12x601	0
469	dpc	1478	7A14x401	0
470	dpc	1478	7A14x501	0
471	dmb	5674	5A12x401	0
7111	dmb	5674	7A14x501	0
...

UNIX-style Account Map:

UID	Charge
1478	7A14x401
3152	5A12x401
5674	5A12x401
...	...

Note that if SITE-style accounting is in use and Account Validation has been enabled, the user-to-account index mappings are already maintained in the Account Validation table. If no additional mappings are needed, there is no need for the site to create a separate site-defined Account Map.

Site-defined account apportionment table

In UNIX-style accounting, the UID (as Account Index) maps only to a specific user. Some sites may wish to apportion different percentages of the charges for a single UID among different project charge codes, but without using SITE-style accounting. HPSS does not provide a means to do this, but a site could implement its own table and utilities to do so. Here is an example of an account apportionment table which a site might implement to do this:

UID	% of (Project(s))
1478	75 (DND) 25 (CBC)
3152	45 (DDI) 25 (DND) 30 (CBC)
5674	100 (DDI)
...

The site would need to write utilities to tie the apportionment table to the HPSS accounting report.

Maintaining site-defined accounting files

Sites which require a site-defined Account Map, apportionment table, or other accounting files must develop tools to maintain and modify the files according to local accounting policies. Tools will be necessary to change, add, or delete an HPSS Account Index and its associated information. Site administrators must implement their own policies on what to do when an account is deleted, a user moves to another project, or a user leaves the system.

Other needed mappings, if any, must be maintained by the site in the Account Map file. When Account Validation is disabled, the site will need to maintain the user-to-account index mappings in the Account Map file as well.

UNIX-style accounting changes of this nature are handled through the normal utilities that set up and modify users and UIDs. A basic set of site-developed utilities are described in more detail below.

Add a user and account

New entries can be made and the next available HPSS Account Index number will be assigned from the free list. The free list will most likely consist of the last assigned number plus one, but could include reclaimed index numbers if a site chooses to reuse Account Index numbers that were previously assigned and no longer referenced. It is not likely that a site will need to reclaim Account Index numbers, but it is an option.

Delete a user

When a user is deleted from the Account Map, the HPSS files must be reassigned to another HPSS Account Index. This should be done from the HPSS client interface side. The **update_flag** should be set to **true** to indicate that this account index number can be reclaimed. The reclaiming tool should check for files using the account number before reclaiming it. When the Account Index is reclaimed, it can be put on the free list to be reused. It is important to keep a copy of the Account Map that corresponds to the HPSS accounting snapshot of storage space in use for that time period so that the proper site information can be matched.

Delete account

When an account is deleted, it is handled in the same manner as when a user is deleted.

Modify account

The entries in the Account Map can be modified to correlate the Account Index to different information, but care should be taken to keep a copy of the corresponding tables for past HPSS accounting runs.

14.2.3.2. Accounting intervals and charges

The time between accounting runs and the charging policy for space usage should be developed after consulting the accounting requirements. The following are some guidelines to consider:

- Accounting should be run at a regular intervals, such as once per month.
- An accounting run may take several minutes, and the storage system will probably be active during the run. The resource usage reported for each user will reflect the resources used by that user at the point when the accounting run encounters that user. This is why accounting represents a blurred snapshot instead of a snapshot at a single point in time.
- Certain accounting information is kept in cache for several minutes after it has changed. For this reason, changes to a user's accounting data may not appear in an accounting report until this period of time has elapsed. Those changes which are still in cache when accounting runs will not appear on the current accounting report, but will appear on the next accounting report.
- The number of file accesses and the amount of data transferred can be taken to represent the activity level of a certain user account in the HPSS system. You may wish to charge specifically for the network and server resources consumed by this activity.
- It may be useful to charge different rates for each Class of Service. For example, a Class of Service that keeps two tape copies of each file will use up more tape cartridges than a Class of Service that keeps only a single copy of each file on tape.

Chapter 15. User interfaces

This chapter describes the configuration information for the user interfaces provided with HPSS for transferring files:

- Client Application Programming Interface (API)
- Parallel File Transfer Protocol (FTP) or PFTP

For other application configuration information, such as HPSSFS-FUSE, consult documentation bundled with the application.

15.1. Client API configuration

The following environment variables can be used to define the Client API configuration. The defaults for these variables are defined in the `hpss_env_defs.h` file.

HPSS_API_AUTHN_MECH

Specifies which authentication mechanism may be used. The values may be `krb5` or `unix`. The default is `HPSS_CLIENT_AUTHN_MECH`.

HPSS_API_BLOCK_SIZE

Specifies the preferred block size for efficient file system I/O. The default is 1048576 bytes (1MiB).

HPSS_API_BUSY_DELAY

Specifies the number of seconds to delay between retry attempts. Note that this value is used both for retrying initialization operations (see `HPSS_API_RETRIES`) and Core Server requests (See `HPSS_API_BUSY_RETRIES`). The default value is "15".

HPSS_API_BUSY_RETRIES

Specifies the number of retries to be performed when a request fails because the Core Server does not currently have an available thread to handle that request. A value of zero indicates that no retries are to be performed (that is, the operation will be performed once), and a value of "-1" indicates that retries should be attempted until either the request succeeds or fails for another reason. The default value is "3" (which will result in up to four attempts).

HPSS_API_DEBUG

Specifies the level of debugging the Client API should produce. By default, no debug logging is produced. For more information on the debugging levels, consult the *HPSS Client API Programmer's Reference*. When API debugging is activated messages will go to `stdout` unless the application sets the log path or the `HPSS_API_DEBUG_PATH` environment variable is set. API debugging can be useful in tracing API problems or understanding network behavior of an HPSS application.

HPSS_API_DEBUG_PATH

Specifies the destination of debug messages; default is `stdout`. If `HPSS_API_DEBUG_PATH` is set to `stdout` or `stderr`, debug messages will be written to the standard output or standard error

I/O streams, respectively. Otherwise, the value will be interpreted as the pathname of a file. The Client API must be compiled with debugging enabled.

HPSS_API_DISABLE_CROSS_REALM

Specifies cross-realm authentication. When cross-realm authentication is disabled, a client will not be allowed to access directories which are located in another security realm. The default value is zero (meaning that cross-cell authentication is enabled). To disable cross realm authentication, HPSS_API_DISABLE_CROSS_REALM must be set to a numeric, nonzero value.

HPSS_API_DISABLE_JUNCTIONS

Specifies whether junction traversal is enabled. When junction traversal is disabled, a client will not be allowed to access directories which require traversal to the directory via a junction. The default value is zero (meaning junction traversal is enabled). To disable junction traversal, HPSS_API_DISABLE_JUNCTIONS must be set to a numeric, nonzero value.

HPSS_API_DMAP_WRITE_UPDATES

Specifies the frequency of cache invalidates that are issued to the DMAP file system while writing to a file that is mirror in HPSS. The value indicates the number of write operations between cache invalidates. The default value is "20".

HPSS_API_HOSTNAME

Specifies the hostname to be used for TCP/IP listen ports created by the Client API. The default value is HPSS_HOST. This value can have a significant impact on data transfer performance for data transfers that are handled by the Client API (that is, those that use the **hpss_Read** and **hpss_Write** interfaces). The default is HPSS_HOST.

HPSS_API_LIMITED_RETRIES

Specifies the number of retry attempts before a limited retry error operation fails. Currently, no operations are defined as "limited retry error operations". The default value is "1".

HPSS_API_MAX_CONN

Specifies the number of connections that are supported by the Client API within a single client process. If HPSS_API_MAX_CONN is set to zero, the number of connections is equal to the default supported by the HPSS connection management software - currently 20. If HPSS_API_MAX_CONN is nonzero, it is the number of connections to be used.

HPSS_API_RETRIES

Specifies the number of retries to attempt when an operation fails. Currently this class of operation includes library initialization and communications failures. A value of zero indicates that no retries are to be performed (that is, the operation will be attempted once), and value of "-1" indicates that the operation will be retried until successful. The default value is "4" (meaning the operation will be attempted up to five times).

HPSS_API_RETRY_STAGE_INP

Specifies whether retries are attempted on attempts to open files in a Class of Service that is configured for background staging on open. A numeric, nonzero value (the default) indicates that opens which would return **-EINPROGRESS** to indicate that the file is being staged will be retried (using HPSS_API_RETRIES and HPSS_API_BUSY_DELAY to control the number and timing of retries), while a value of zero indicates that the **-EINPROGRESS** return code will be returned to the client. The default value is "1".

HPSS_API_REUSE_CONNECTIONS

Specifies whether TCP/IP connections are to be left open as long as a file is open or are to be closed after each read or write request. A numeric, nonzero value will cause connections to remain open, while a value of zero will cause connections to be closed after each file access. The default value is zero.

HPSS_API_RPC_PROT_LEVEL

Specifies the RPC protection level. Valid values are `connect`, `packet`, `packet integrity`, and `packet privacy`. The default is `HPSS_RPC_PROT_LEVEL`.

HPSS_API_SAN3P

Specifies whether to enable SAN3P transfers. Valid values are "on" and "off". The default is on.

HPSS_API_SITE_NAME

Specifies the site name to be used when initializing the HPSS security services. The default value is `HPSS_SITE_NAME`.

HPSS_API_TOTAL_DELAY

Specifies the number of seconds to continue retrying requests. A value of zero indicates that no there is no time limit. The default value is zero.

HPSS_API_USE_PORT_RANGE

Specifies whether the HPSS Movers should use the configured port range when making TCP/IP connections for read and write requests. A numeric, nonzero value will cause Movers to use the port range, while a value of zero will cause Movers to allow the operating system to select the port number. The default value is zero.

HPSS_KRB5_KEYTAB_FILE

Specifies the name of the file containing the security keys necessary for successfully initializing the Client API for Kerberos authentication. The default is `auth_keytab:/var/hpss/etc/hpss.keytab`.

HPSS_UNIX_KEYTAB_FILE

Specifies the name of the file containing the security keys necessary for successfully initializing the Client API for UNIX authentication. The default is `auth_keytab:/var/hpss/etc/hpss.unix.keytab`.

HPSS_SEC_MUTUAL_AUTH

Specifies whether clients should utilize mutual authentication security. This forces the client to be authenticated by the server (which always happens), and the server's credentials authenticated by the client. This can prevent certain man-in-the-middle attacks. The default is `true` (on). Due to current security subsystem limitations (such as in UNIX system authentication), mutual authentication may be turned off by setting this variable to `off`. There is some amount of overhead to mutual authentication for first-time client startup, which may be another reason a site would disable it.

15.2. FTP/PFTP daemon configuration

The `pftp_client` binary is an enhanced FTP client providing parallel data transfer facilities, HPSS-specific features, and transfer of files greater than 4 GB. The HPSS binary, `hpss_pftpd`, is an enhanced FTP Server facilitating parallel transfers into HPSS. This binary *only* functions with HPSS. If you are

using an "ftp" binary or a "pftp_client" binary to access HPSS the site *must* run the hpss_pftpd binary. The hpss_pftpd binary provides FTP services for both the standard system FTP client and the parallel-enhanced HPSS pftp_client.

Note: Not every feature of every system FTP client will be honored by the hpss_pftpd. Additionally, system FTP clients cannot take advantage of the parallel transfer features nor the increased file size capacity provided by the hpss_pftpd. If a site wishes to provide FTP services to both the system FTP server, ftpd, and HPSS, hpss_pftpd, it will be necessary to set up the two services to listen on different ports.

There are four steps that must be performed prior to using the FTP user interface to transfer HPSS data:

1. Verifying the FTP initial configuration
2. Configuring the FTP daemon syslog
3. Defining the FTP access
4. Creating FTP users

These steps are described in more detail in the paragraphs that follow.

Step 1. Verifying the FTP configuration

During the HPSS infrastructure configuration phase, the following files were created or modified by the **mkhpss** script:

- In the `/etc/services` file, verify that the `hpssftp` and `hpssftp-data` entries were added correctly. Ports 20 and 21 are the ports of choice to prevent having to educate every user to alternate ports. By default, **mkhpss** uses ports 4020/4021. It may be necessary to move the existing `ftp` and `ftp-data` to other ports when assigning the `hpssftp` and `hpssftp-data` to ports 20/21.
- In the `/etc/inetd.conf` or `/etc/xinetd.conf` file: Verify that the **hpssftp** entry was added correctly. The only flag that has relevance is the **-cFileName** flag which specifies an alternative `HPSS.conf` file (located in `/var/hpss/etc`).

Refresh **inetd** or restart **xinetd** after making these changes:

```
% refresh -s inetd          # (AIX only)
% /etc/init.d/xinetd restart # (Linux only)
```

A "TCP Wrapper" application may be used for initiating the HPSS PFTP Daemon to make it more secure; however, such a program is not provided by the HPSS project.

The file `/var/hpss/etc/HPSS.conf` provides the options for both an HPSS PFTP daemon and a non-HPSS PFTP daemon (DISCOM PFTP daemon). This file should be customized as needed. Refer to the `HPSS.conf` man page or the *HPSS.conf configuration file* appendix of the *HPSS Installation Guide* for details.



It may be necessary for a site to merge older copies of the `HPSS.conf` file with the template if modifications have been made since the prior release. There is no conversion job to perform this task.

Step 2. Configuring the FTP/PFTP daemon syslog

The PFTP daemon attempts to write to the system log (using **syslog()**). To enable this output, follow local system procedures for capturing this information. The PFTP daemon writes to the **LOG_DAEMON** syslog facility.

Step 3. Defining the PFTP access

In the following statements, **{FTPBaseDir}** defaults to `"/var/hpss"` and **{ftpass}** defaults to `"ftpass"`. These values may be changed in the `HPSS.conf` file. The **{FTPBaseDir}** `/etc/{ftpass}` file defines access, options, and restrictions for the server. An `ftpass.tpl` file is provided in the `$HPSS_ROOT/config/templates` directory for the HPSS administrator to customize appropriately. The pound sign (`#`) at the beginning of a line indicates a comment. Options are enabled by removing the comment symbol (`#`) and conversely disabled by introducing this symbol. Some options are specific to individual machines, so each HPSS PFTP daemon should read its own **{ftpass}** file.

The access options can be defined as follows:

```
# Define a class of users.
# class <class> { anonymous | guest | real } <user> [ <user> ... ]
#
# <user> is a full network address, including wildcarding, e.g.,
# *.nsl.nersc.gov.
class hpss_class real *
class hpss_anon anonymous *
#
# Define a limit on the number of users from a certain class
# that can be using the ftp server at specific times.
# limit <class> <n> [<times>] [<msg_file>]
#
# The <times> clause has the following format:
# <day><time>--<time>|<day><time>--<time>|... (no white space!)
# or "Any"
# where <day> is one of "Su", "Mo", "Tu", "We", "Th", "Fr",
# "Sa", "Wk"
# "Wk" matches any day of the week
# <time> is the time in 24-hour format
# "Any" matches any day and time
# (used to specify <msg_file> for all days and times)
# limit hpss_class 1 Tu0600-1200
limit hpss_class 3
#
# Display a banner prior to the user prompt (before log in.) Very
# useful for informing the user that he/she is dealing with the
# HPSS file systems rather than the local file system of the
# host executing the FTP Daemon. This is highly recommended!

banner <pathname/filename>

# Control for logging (sent to syslog()).
log [ commands ] [ anonymous ] [ { inbound } ]
[ transfers ] [ guest ] [ { outbound } ] [ real ] [debug]

# Set the maximum number of login attempts before closing the
# connection:
loginfails <login_failure_threshold>

# Determine the appropriate behavior when needed data must be staged.
```

```

# <stage-code> can have several values:
# -1 wait forever
# 0 do not wait (default) - inform user that data needs staging and
# exit
# >0 wait for n seconds - if staging takes more than n seconds, inform
# the user and exit
# note: This option requires that the relevant storage class be
# configured in a compatible way, or behavior may not be as
# expected. For example, setting wait_on_stage to -1 (wait forever)
# for a SC with no-stage set would make little sense.
wait_on_stage <stage-code>
#
# Define users as being guests (which means they cannot
# cd out of their own subdirectory).
guestgroup <group> [ <group> ... ]

# Deny access to a set of users, <addrglob> is full network address,
# including wildcarding, e.g., *.nsl.nersc.gov.
deny <addrglob> [ <msg_file> ]

# Send shutdown message to current FTP users.
shutdown <shutdown_info_pathname>

# Send message to users, possible on access to a directory.
message <pathname> [ <when> ]

# Display prompt to read a file, possible on access to a directory.
readme <pathname> [ <when> ]
# The <when> clauses may be either "login" (in which case the file
# is displayed when the user logs in) - Applicable only for anonymous
# or "cwd=[<directory_name> | *]",
# (in which case the file is displayed when the user changes to the
# named directory, with '*' wildcarding all directories).
# ONLY displays message on FIRST entrance. Directory MUST be full path

```

- **Message/readme/banner/shutdown** (message lines) are text files, with the following keywords (all one character in length) recognized, identified by a preceding %:

Table 15.1. Banner keywords

Keyword	Description
M	Class limit
T	Current time on FTP Daemon
C	Current working directory
R	Remote hostname
L	Local hostname
U	User name
s	Shutdown time
d	User disconnect time
r	Connection deny time
%	%

- The format of the **<shutdown_info_pathname>** file is:

```
<year> <mon> <day> <hour> <min> <deny> <disc>
```

Message lines contain keywords mentioned above. For example:

```
1994 1 28 16 0 120 30
System shutdown at %s (current time is %T)
New FTP sessions denied at %r
FTP users disconnected at %d
```

indicates that shutdown will be 1/28/94 at 16:00, with users disconnected at 15:30 and sessions denied at 14:40 (the 120 indicates 1 hour, 20 minutes).

- The **<when>** clauses may be either **login** (in which case the file is displayed when the user logs in) or **cwd=[<directory_name> | *]** (in which case the file is displayed when the user changes to the named directory, with an asterisk (*) wildcarding all directories).
- The **<times>** clause has the following format:

```
<day><times>-<time> | <day><time>-<time> | . . .
```

where **<day>** is one of **Su, Mo, Tu, We, Th, Fr, Sa, Su, Wk**, or **Any** and **<time>** is the time of day on a 24-hour clock. For example, to specify Tuesday between 8:00 AM and 4:00 PM:

Tu0800-1600

The **Wk** entry for the **<day>** clause indicates a weekday (Monday through Friday), and **Any** indicates all days of the week.

Step 4. Creating FTP users

In order for an HPSS user to use FTP, a UNIX or Kerberos (or both) userid and password must be created. Refer to *Section 3.3.2.1, "The **hpssuser** utility"* for information on how to use the **hpssuser** utility to create the userid and password and set up the necessary configuration for the user to use FTP. Note that this step should not be done until the Core Server is running so that the **hpssuser** utility can create the home directory for the FTP user.

The **/opt/hpss/bin/hpss_pftppw** utility can be used to set the encrypted passwords in the **/var/hpss/etc/passwd** file. The syntax for this utility is as follows:

```
hpss_pftppw <userid> [<password file pathname>]
```

The utility will prompt the user for the old and new passwords. If root has to change the password it is necessary to manually edit the file and remove the existing encrypted password. The password file pathname argument can be used to specify a password file other than the default file, **/var/hpss/etc/passwd**.

To enable anonymous FTP, the **"hpssftp"** user must be defined in either the HPSS FTP password file or in the Kerberos KDC and LDAP registry (depending on which authentication mechanism is enabled). In addition, the entry for the **"hpssftp"** user must contain a home directory defined to be a non-NULL value. Anonymous FTP users will have the "root" directory set to **"/anonymous"** to prevent anonymous users from accessing other portions of HPSS. This means that for anonymous access the directory **/anonymous** *must* exist.

To disable anonymous FTP, do either or both of the following:

- Remove the `hpss_anon` class from the `ftpaccess` file
- Add **hpss_ftp**, **anonymous** or **guest** to the HPSS FTP user file (normally "`/var/hpss/etc/ftpusers`").

Security policies and the security requirements should be carefully examined and thoroughly understood before allowing "anonymous" FTP.

Chapter 16. Backup and recovery

This chapter discusses the following operations:

- Backup and recover HPSS metadata
- Backup HPSS environment
- Recover HPSS user data
- Handling DB2 space shortage

16.1. HPSS metadata backup and recovery

Each HPSS site is responsible for implementing and executing the HPSS metadata backup and recovery process. The HPSS administrator must ensure that HPSS metadata is backed up on a regular basis. In addition, the HPSS administrator must verify that the backup files can be used to restore the HPSS metadata up to the point of failure. Sites can choose to use a commercially available backup software tool, such as the Tivoli Storage Manager (TSM), to backup the HPSS metadata. Optionally, a set of simple and customizable backup scripts can be developed to backup the HPSS metadata. Contact HPSS support to obtain sample backup scripts and guidance for implementing the DB2 backup process for your site.

The following sections provide an overview of the metadata backup and recovery process.



The backup topics provided in the following sections are meant only as an overview. The DB2 Data Recovery and High Availability Guide and Reference should be consulted in order to obtain a complete understanding and explanation of issues and procedures associated with enabling full recoverability of DB2 data.

16.1.1. HPSS administrator responsibilities for DB2

While DB2 has sophisticated backup and recovery abilities, it is fundamentally the responsibility of the HPSS administrator to ensure that these facilities are configured and working correctly. This section outlines some of those responsibilities.

While storing the DB2 data tables on reliable disk storage systems can be very useful if it becomes necessary to recover the database, it is essential that the database be backed up regularly, and that the backup media be reliable. Many sites keep multiple copies of the backup media in separate locations. It is also extremely important that site administrators verify that they can restore a failed database from the backup media. Time should be scheduled on an occasional basis to perform practice recoveries from the backup media. This could be done by creating a separate instance of the HPSS DB2 system on a separate computer and restoring it from the primary instance backup media.

HPSS databases should be backed up regularly. Most sites back up the entire database once each night, using the online backup procedure. Note that Online backups are only available for database configurations that enable log archiving. (see *Section 16.1.2.1, “Configuring DB2 for online*

backup”). Online backups do not interfere with normal HPSS processing. It is equally important to monitor the online backup process and verify it is completing normally.

As HPSS operations take place, DB2 writes transaction log entries describing the changes that have been made to the database. These logs allow DB2 to recover all changes made to the database since the time of the last database backup, forward to the time when DB2 was stopped by a crash, hardware failure, power outage or whatever. It is vital that the DB2 log files be hosted on a highly reliable disk system. Furthermore, DB2 log mirroring must be used to protect this information on two separate disk systems. The disk systems must also be protected using RAID 1 or RAID 5 configurations. Reliable access to the log files is more important than performance. Loss of the DB2 log files can result in a significant impact to the operation of the system, extensive downtime to repair, and potential loss of end-user data. Proper protection of the DB2 log files and associated archived backups is critical to a site's HPSS operations.

Similarly, the HPSS administrator must regularly verify that completed DB2 log files are being stored on reliable media until they are no longer needed. Many sites make multiple copies of DB2 log files and store them in separate physical locations.

In case of a catastrophic failure of the HPSS database host computer, the database can be restored to a consistent state at the point in time of the failure, if the administrator has a reliable series of backup images and log files. After the failure that caused the database failure is corrected, the HPSS databases can be restored using the backup files, and then by replaying the DB2 log files. Much of this process happens automatically when DB2 is restarted, but if you have any questions about these requirements and procedures, contact HPSS support.

Another responsibility of the HPSS administrator is to periodically perform RUNSTATS operations on the database tables. RUNSTATS is a DB2 process that analyzes the tables and produces statistics that are vital to DB2's performance. Without good statistics, DB2 may create very inefficient database access plans that result in very poor performance. Although frequent RUNSTATS are not necessary on HPSS databases, the administrator should have a plan to perform them on some sort of regular schedule consistent with local requirements. You should discuss this with HPSS support.

It is also important to have a plan to perform REORGCHK operations from time to time to monitor the organization of the database tables. As a consequence of normal use, a database table may become significantly disorganized. Database performance can be enhanced by periodic reorganizations, some of which can be performed with HPSS in operation. You should also discuss this with HPSS support.

16.1.2. Overview of the DB2 backup process

Without proper and systematic metadata backup, a database failure could result in total loss of the HPSS system. For this reason, backing up the DB2 databases containing HPSS metadata is essential.

DB2 databases can be backed up in one of two ways:

1. Online backups are made that can be restored to the point of failure by "rolling forward" archived logs of transaction activity since the backup.

This is the recommended method for making HPSS backups.

2. An offline snapshot of the database data is taken at a particular point in time.

Offline backups provide recovery only to the point in time at which they were made, and they require that there be no other activity in the database while they are being made. *Therefore, this method is not recommended for use with HPSS.*

Since HPSS installations require the use of online backups, it is useful to understand the types of backup and restore operations that can be performed with this backup type. Online backups can be described as "full", "incremental", or "delta". Full backups record the full contents of the database at a point in time. Incremental backups record all the changed data since the last full backup. Delta backups record only the data that has changed since the most recent backup (of any type). Obviously, a full backup is necessary as a starting point in the recovery of a database. Incremental and delta backups help reduce the number of transaction logs that must be retrieved and applied during a restore operation.

Consider a site that keeps the following schedule for backing up a database:

```
Sunday: full
Monday: delta
Tuesday: delta
Wednesday: incremental
Thursday: delta
Friday: delta
Saturday: delta
```

Under this scheme, if the database fails between the Thursday and Friday backups, only a few backup images would be needed for a restore: Sunday's full backup is needed as a starting point. Wednesday's incremental backup includes all changes since Sunday, so it makes retrieval of Monday's and Tuesday's backups unnecessary. Beyond that, Thursday's delta backup is needed, as well as any transaction logs archived after Thursday's backup. So the needed images are: Sunday, Wednesday, Thursday, and only those logs between Thursday's backup and the point of failure.

The way DB2's backup and restore commands are implemented makes restoration simple. Basically, only the latest backup image needs to be specified. DB2 records the backup history in each backup, so it can determine which images and logs are necessary. Assuming an automated method is used for backup storage and retrieval, it can also fetch the images and restore from them. See the *DB2 Data Recovery Guide* for details.

16.1.2.1. Configuring DB2 for online backup

Several steps must be taken to prepare a DB2 database for online backup. Initially, a newly-created database is configured to perform offline backups. The database's configuration must be updated to enable log archiving. Online backups are only available when log archiving is enabled. An initial, offline backup must then be made as a starting point for the backup process. The database will be inoperative until this initial full backup is completed. An outline of the major steps to set up a database for online backup is:

1. Configure backup software and make certain data can be archived into it and retrieved from it.
2. Update the database configuration to enable log archiving, this will place the database in "backup pending" state.
3. Perform initial full, offline backup (database returns to normal state).

4. Schedule subsequent backups using **cron**.

To back up a DB2 database, invoke the **db2 backup** command using the DB2 command line processor. Typically, the administrator would put this command into a script along with any necessary parameters, such as the type of backup, the database partitions, the destination, and other important information required by DB2 to conduct its backup. The script would then be invoked from **cron**.

For details on configuring DB2 for backup, administrators should consult the *DB2 Data Recovery Guide* as well as the documentation for their backup software.

16.1.3. Overview of the DB2 recovery process

Upon discovering a damaged DB2 container, you must first determine the level of hardware recovery and problem determination available to you. For example, whether or not you were utilizing some level of RAID and can recover a failed disk from a good disk. The second step in recovering damaged or inaccessible DB2 containers is considering any software recovery tools available. For example, if you were utilizing backup software, such as TSM, this would enable you to replace the bad hardware and recover onto the new hardware device to a particular point in time. For details on recovering a DB2 database, see the IBM redbook titled "Data Recovery and High Availability Guide and Reference" for the appropriate version of DB2 installed on your system. This guide will help you decide how best to recover your DB2 database or damaged container.

16.2. HPSS system environmental backup

The HPSS administrator must also ensure that the HPSS system environment is backed up on a regular basis and to verify that the backups can be used to restore the environment. The following sections describe the environmental data that must be safeguarded.

16.2.1. HPSS file system backup

The following HPSS file systems should be backed up using the site backup procedure:

- /opt/hpss
- /var/hpss
- /var/hpss/adm/core
- {HPSS metadata backup path}/cfg
- {HPSS metadata backup path}/subsys1
- /db2/backups/ldap (if LDAP is used)
- /var/hpss/hpssdb
- /var/hpss/hpssdb/ldap (if LDAP is used)
- {DB2 log path}/cfg
- {DB2 log path}/subsys1

- {DB2 log path}/ldap (if LDAP is used)
- {DB2 log archive path}/cfg
- {DB2 log archive path}/subsys1
- {DB2 mirror log path}/cfg
- {DB2 mirror log path}/subsys1
- {DB2 mirror log archive path}/cfg
- {DB2 mirror log archive path}/subsys1
- {HPSS secondary metadata backup path}/cfg
- {HPSS secondary metadata backup path}/subsys1
- /<metadata backup file system>
- Other site-specific file systems

16.2.2. Operating system backup

It is also necessary to perform appropriate OS backups on a regular basis. However, these activities are outside the scope of this document. Refer to the relevant OS documentation for the appropriate backup procedures.

16.2.3. Kerberos backup

If Kerberos authentication is configured, periodic backups of the Kerberos infrastructure will need to be made. If **mkhps** was used to configure Kerberos, the Kerberos components that must be backed up are placed in the following locations:

- The Kerberos client configuration file is `/etc/krb5.conf`.
- The default keytab file is `/etc/v5srvtab`.
- If a Key Distribution Center (KDC) is running on a node, all files related to it are placed in a single directory: `/var/hpss/krb5kdc`.

If you configured Kerberos manually, you'll need to determine the file locations yourself. The files that should be backed up are:

- Client configuration file
- KDC configuration file
- KDC `acl_file`
- KDC `admin_keytab`
- KDC `dict_file`

- KDC database_name
- KDC key_stash_file
- KDC default_keytab_name

16.3. HPSS user data recovery: Recovering HPSS files from damaged HPSS volumes

HPSS tape or disk volumes can become unreadable in cases of physical damage to the media or of being accidentally over-written by a non-HPSS application. If an HPSS tape or disk volume is damaged such that HPSS cannot retrieve the user data stored there, its data might be recoverable using the HPSS **repack** or **recover** utility.

To determine whether recovery is needed, use this procedure:

1. Determine the name of the potentially damaged volume.

When HPSS cannot read a disk or tape volume, the HPSS Mover will issue alarm messages which will be displayed in the SSM Alarm and Events window. The alarm messages will contain the name of the physical volume for which the error occurred. Record the volume name.

2. Determine if the volume is actually damaged.

Check the central HPSS log for the alarm message and any related messages for the source of the problem. Some problems are configuration issues. In particular, for errors for tape volumes, the system administrator should check the block size of the tape device on which the error occurred. HPSS tape devices should be configured with variable block sizes (see *Section 7.1.2, “Enable variable block sizes for tape devices”*).

If the problems are not configuration related and the volume does appear to be damaged, **repack** or **recover** the suspect volume. See the **repack** and **recover** man pages for instructions on running these utilities from the command line.

Repack reads the data from the damaged volume and copies it to a good volume at the damaged level of the hierarchy. **Recover** reads the data from an alternate copy on a volume at another level of the hierarchy and copies it to a good volume at the damaged level of the hierarchy. **Recover** works by staging the data from the alternate copy to level 0 of the hierarchy and forcing the migration of the data all the way down the hierarchy. **Recover** can also be used to clean up the metadata for the damaged volume if the data cannot be restored.

In some situations, one utility may be better suited than the other. In other situations, a combination of the utilities may be best. Some considerations:

- **Recover** requires an alternate copy, which is another complete and permanent copy of the data at another level of the hierarchy. A permanent copy is a copy that is not subject to purge. The system must be configured in advance to provide two or more alternate copies. If there are no alternate

copies of the data, the data cannot be restored by the **recover** utility. See the **recover** man page for more information about alternate copies. **Repack** does not require an alternate copy.

- **Repack** must be able to read the damaged volume. If the volume is unreadable, the data cannot be restored by the **repack** utility. In some cases, **repack** is able to read part but not all of the data and restore the portions it can read. **Recover** does not read the damaged volume.
- **Recover** will re-migrate the data to all levels of the hierarchy. **Repack** will copy the data only to the required level.
- Since **recover** stages the data from the alternate copy to level 0 of the hierarchy, there must be sufficient space at level 0 to hold all the alternate file copies. If there is not enough space, the files still might be restored by executing **recover** in phases. **Recover** can be started and allowed to recover as many files as it is able given the space limitations. Then, once space is freed up at level 0 by normal purge processing, **recover** can be restarted in "--retry-failures" mode to retry the remaining files. It can be restarted multiple times if necessary until all files are recovered. **Repack** does not depend upon staging the file to level 0 or the amount of space there.
- For tape volumes at any level to which data is migrated from disk, **recover** by default aggregates the recovered files on the new volume. If for any reason the recovered files cannot be aggregated, **recover** attempts to place them on the new volume as single files. **Recover** does not aggregate the recovered files if the tape volume is at level 0 or if data is migrated to the volume from tape. **Repack** aggregates files to tape by default, but can be executed with an option not to aggregate. The behavior of **repack** is the same for tape at any level, regardless of where the tape is in the hierarchy.

If all of the data cannot be restored to another volume by either utility, the **recover** utility must be used to clean up the metadata for the remaining data so that the volume can be removed from the system.

A good general approach is:

- If you have alternate copies:
 1. Run **recover** first to try to restore the data from an alternate copy.
 2. If you can't get all the data back from **recover**, run **repack** to try to get the rest of the data back directly from the damaged volume.
 3. If you still can't get all the data back, run **recover** again to clean up the metadata for the unrecoverable data. This is a last resort and all other options of recovering the data, including running **recover** with the **--retry-failures** option, should be attempted first.
- If you don't have alternate copies:
 1. Run **repack** to try to get as much of the data as possible back directly from the damaged volume.
 2. If you can't get all the data back, run **recover** to clean up the metadata for the unrecoverable data. This is a last resort and all other options of recovering the data, including running **recover** with the **--retry-failures** option, should be attempted first.

Once **recover** is executed on a volume, it must either run to completion or be cancelled. HPSS will not allow the volume to be removed from the system using the "delete resources" operation so long as

there is any outstanding recover operation in progress. Even if the recover utility is no longer running and the volume is empty, HPSS will consider the recover operation to be still in progress until it is marked as complete or cancelled. If **recover** is run for a volume and does not recover all of the files, and **repack** is then run and does finish retrieving all of the data, the recover operation must be cancelled. See the **recover** man page for instructions on cancelling the recover operation.



If the first repack attempt fails, it may be worthwhile to retry the repack several times. This is especially true if alternate copies are not available for the storage level that contains the suspect volumes. There are a number of intermittent or environmental errors that may cause an attempt to repack a volume (especially a tape volume) to fail, while a subsequent attempt may succeed (for example, a failing drive, problems reading a volume in a particular drive, or a drive that requires cleaning). If the suspect volume is a tape, try to get it mounted on a number of drives during the repack effort by locking the drives in which it has problems. Sometimes a tape will fail on one drive but can be read on another.

16.4. DB2 monitoring

For the convenience of HPSS administrators, HPSS monitors several critical aspects of the underlying DB2 databases. The amount of free space in the global (often called "cfg") database is monitored by the Core Server running the lowest numbered subsystem. Similarly, the amount of free space in each of the subsystem databases (usually called subsysx, where "x" is the subsystem number) is monitored by the Core Server that runs that subsystem. The configuration settings for these monitors are found in the global and subsystem configurations. The interval at which the monitors run and the thresholds at which they send alarms can be configured to suit the HPSS administrators' needs.

Starting with HPSS version 6.2.2, a DB2 Log Monitor is available in HPSS. This monitor periodically checks the DB2 transaction log files to make sure the primary and mirrored logs are congruent and apparently performing normally. It also scans the DB2 diagnostic log, `db2diag.log`, for indications of any sort of problem with the transaction logs. If any errors are reported in `db2diag.log`, the HPSS DB2 Log Monitor sends an alarm and indicates the file path and line number where the error was reported. The administrator should be able to open `db2diag.log` and examine the problem report in detail. The error is reported one time for each occurrence in `db2diag.log`. This monitoring function does not interfere with normal maintenance of `db2diag.log`. When the file is trimmed back, or removed, the monitor automatically tracks these changes.

16.5. DB2 space shortage

As HPSS adds files to the database, the DB2 container spaces will be consumed. If the DB2 Space Monitor is configured to monitor free space, an alarm will be generated when free space falls below a set threshold. The following sections discuss how to handle the appropriate space shortage issue.

Before performing any of the suggestions or steps listed in this section, consult the appropriate *DB2 Administration Guide: Implementation* for the appropriate version of DB2. It would also be wise to verify any significant changes to, or reallocation of, production DB2 resources with HPSS support.

16.5.1. DMS table spaces

Capacity for a DMS table space is the total size of containers allocated to the table space. When a DMS table space reaches capacity (depending on the usage of the table space, 90% is a possible

threshold), you should add more space to it. The database manager will automatically rebalance the tables in the DMS table space across all available containers. During rebalancing, data in the table space remains accessible.

A DB2 container may contain more than one physical disk, but a physical disk can be used in only one DB2 container. We recommend allocating the entire physical disk to a DB2 container and not leaving over any free space on the physical disk since it will be considered wasted space once the disk is allocated to the DB2 container.

The *DB2 Administration Guide: Implementation* states that you can increase the size of a DMS table space by adding one or more containers to the table space.

To add a container to a DMS table space using the Control Center:

1. Expand the object tree until you see the Table Spaces folder.
2. Right-click the table space you wish to add the container to, and select Alter from the pop-up menu.
3. Select the Containers tab, click Add, complete the information, and click OK.
4. Click Ok.

To add a container to a DMS table space using the command line, enter:

```
ALTER TABLESPACE <name> ADD (DEVICE '<path>' <size>)
```

For example:

```
ALTER TABLESPACE CFG ADD (DEVICE '/dev/rhd5' 10000)
```

Note that size is in pages based on the page size of the existing table space you are adding the container too. In this case, container refers to the DEVICE.

Note that the ALTER TABLESPACE statement allows you to change other properties of the table space that can affect performance. Refer to "Table Space Impact on Query Optimization" in the *DB2 Administration Guide: Performance* for more information.

16.5.2. SMS table spaces

In general, you cannot extend the size of an SMS table space very easily because SMS capacity depends on the space available in the file system and the maximum size of the file supported by the operating system. You should be able to increase the size of a file system using the normal operating system facilities for increasing file system space.

Chapter 17. Management tools

17.1. Utility overview

HPSS provides a variety of command-line utility programs to aid in the use and management of the system. The programs can be grouped into the following major categories:

17.1.1. Fileset and junction management

These programs work for HPSS filesets.

crtjunction

Creates an HPSS junction.

deljunction

Deletes an HPSS junction.

lsfilesets

Lists all HPSS filesets.

lsjunctions

Lists all HPSS junctions.

17.1.2. Tape management

These programs manage the life cycle of tapes in the system.

repack

Consolidates data on disks and tapes by moving storage segments from sparsely populated volumes to other volumes. The emptied volumes are then available for reuse or removal from the system. This is the principal tool for recovering data from degraded RAIT volumes and writing it to new volumes.

reclaim

Reinitializes an empty tape for reuse after repack has emptied it, or after it has become empty through attrition or repacking. Reclaim will not reinitialize a tape volume that has been retired, or degraded RAIT volumes.

retire

Marks a disk or tape volume as read-only and not reclaimable.

remove

Deletes empty disk and tape volumes from the Core Server's metadata, then optionally exports the volumes from the PVL/PVR.

recover

Salvages data from damaged volumes by reading from other copies or other levels in the hierarchy.

lbp_verify

Mounts tape volumes and verifies the data against stored CRCs.

These tools deal with shelving tapes:

shelf_tape

Selects and moves tapes from a tape library to an offline storage area. Additionally, shelf_tape can check in tapes that previously existed on shelf.

shelf_tape_util

Performs various administrative functions related to shelved tapes.

17.1.3. System info

dump_sspvs

Lists the physical volumes known to a particular subsystem, including statistics about file aggregates on tapes and RAIT volume status.

dumpbf

Lists the storage characteristics of a particular bitfile.

lsvol

Lists the path names of bitfiles that are located on a disk or tape volume.

dumpppv_pvl

List the physical volumes managed by the PVL.

dumpppv_pvr

List the physical volumes managed by a particular PVR.

lshpss

Lists various views of the system configuration.

lsrb

Lists a list of bitfiles that have not been accessed since a given date/time.

lspvhist

List physical—as opposed to virtual—tape volume history information.

mpps_reporter

Produces human-readable output from MPS summary files.

plu

Lists files that have a purge record in HPSS metadata. Can be used to see which bitfiles have been purge locked (for example, bitfiles that cannot be purged).

lscos

Lists all bitfile COS change records by subsystem. It also allows an administrator to cancel all COS change requests for a particular subsystem, and print a summary of COS change records by stream.

device_scan

Lists information all SCSI-attached devices including media-changers, disks, tape.

scsi_home

Displays SCSI home locations.

17.1.4. System management

hpssadm

Allows administrators to perform some administration of HPSS without the graphical user interface. Not all SSM functionality is available. But it enables the administrator to perform all configuration operations, to display alarms and events, to display all the lists provided by the SSM GUI, and to manage HPSS servers, devices, storage classes, PVL jobs, volumes, filesets, and junctions. It may be run interactively or in batch mode.

hpssgui

Starts the Graphical User Interface program which is the primary administrative interface for HPSS.

hacl

Allows viewing and editing of HPSS access control lists (ACLs).

hpssuser

Adds, manages, and removes HPSS users.

rtmu

(Real-Time Monitoring Utility) Displays the state of requests in the HPSS system.

hpss_avaedit

Allows HPSS administrators at sites using SITE-style accounting and account validation to edit the relationships between user IDs, account IDs, and account names.

hpss_server_acl

Updates authorization data in the database.

hpss_krb5_keytab

Generates a keytab in the form usable by the **hpssadm** program.

hpss_krb5_rmxcred

Removes expired cached Kerberos credentials from the default credential directory.

hpss_unix_group

Manages HPSS UNIX group file.

hpss_unix_keygen

Create an HPSS UNIX master key.

hpss_unix_keytab

A keytab generation program for the HPSS UNIX native security service.

hpss_unix_passwd

Changes the HPSS UNIX native security service password for an HPSS user.

hpss_unix_user

Manages HPSS UNIX password file.

hpss_managetables

Creates, deletes and modifies databases, database tablespaces, tables, views and constraints. It should be used by hand only by extremely knowledgeable administrators. This program can destroy database tables.

xmladdindex

Creates an XML index.

xmldeleteindex

Deletes an XML index.

schema_registrar

Registers an XML schema with the database.

hpsssum

Checksums HPSS or local files

17.1.5. User interfaces

pftp_client

Used to perform parallel file transfers between a local and remote host.

17.1.6. Testing and debugging

scrub

A general-purpose HPSS test shell; a driver for the Client API library. It can perform various data operations, such as reading and writing undifferentiated data, migrating and purging specific files, and browsing the HPSS name space. **Scrub** also supports doing some simple operations with user-defined attributes. It does not transfer external files into or out of HPSS.



Scrub is a test tool used by developers for testing the Client API, Core Server, and other areas of HPSS. As such, it does not provide many of the protections offered by other interfaces into the system. For this reason, it must be treated by production sites with caution. It is available for limited use by production sites for information gathering and for other operations described in official HPSS documentation. To prevent unintended consequences, any other use at production sites should be cleared with HPSS support.

17.1.7. Unsupported tools

Tools in this category are not formally supported. They are provided as-is and are found in the `tools/unsupported` portion of the HPSS source tree.

disk_allocation_stat

Calculates and displays statistics on disk usage and fragmentation for disk storage classes.

ftp

(DISCOM PFTP Daemon) This is a portable version of a standard UNIX FTP daemon which supports the HPSS Parallel FTP protocols. By replacing the System FTP Daemon and configuring the DISCOM Daemon, significant improvements can be made in moving data between two UNIX-based systems. Both a Kerberized version and a non-Kerberized version are available for a number of hardware/software platforms. The HPSS Parallel FTP Client should be used to talk to this server. Do not confuse this with the HPSS FTP daemon.

hpsstats

Reports HPSS statistics.

Load tools

These utilities reside in `tools/unsupported/src/load`. Each tool allows the manual loading, editing and deletion of one or more metadata records.

magic

Allows direct manipulation of name space objects.

nsde

Displays and edits Name Service database tables.

PVrr

Displays the number of PV retrieval requests that have been made on a volume.

removesegs

Searches for and removes disk and tape storage segments that are not pointed to by any bitfile metadata. It requires at least two separate searches of the bitfile metadata over a period of time before it will remove orphaned segments. It can also be used to reset the number of segments counter in disk and tape storage maps and to repair certain BFS metadata defects.

showdiskmaps

Sends a command to the Core Server in the selected storage subsystem to dump its in-memory disk space allocation maps to a file in `${HPSS_PATH_TMP}`, then displays that information on standard output. The tool's output can give a point-in-time snapshot of how space is allocated across the disk VVs in a system and how busy those disk VVs (and the movers that operate them) are. The output includes information about free, usable, and actual space; number of I/O operations in flight; and the number and nature of blocks available for allocation.

17.2. Tape queue monitoring

HPSS supports dumping out tape queue data. When read, write, dismount, or copy requests are made to tape resources, they are queued up and ordered using the Core Server's configured ordering method.

Periodically, the entire tape queue layer's status is dumped to local disk in order to facilitate administrative queue monitoring. The location of this dump is `$HPSS_PATH_TMP/tape-schedule-
<subsystem id>.json` (for example, `/var/hpss/tmp/tape-schedule-1.json`). The file is in JSON format and is overwritten every five seconds (by default). The dump interval can be modified by setting the environment variable `HPSS_SCHEDULE_DUMP_INTERVAL_SECONDS` and restarting

the Core Server. It is not recommended to decrease this interval below five seconds. This dump remains enabled so long as the path can be opened for writing. These fields are subject to change between HPSS releases.

17.2.1. Example

```
{
  "VV": [
    {
      "Queue": {
        "Status": "Active",
        "Sessions": [
          {
            "Activity Id": 0,
            "Start Section": 28,
            "Is Complete": false,
            "TimeOrder": 0,
            "Score": 0,
            "Request Id": 1003421707,
            "Is Reader": true,
            "Status": "RUNNING",
            "Start Offset": 0,
            "Wait Time": 95745,
            "Data Length": 1024,
            "Is Dismount": false,
            "Is Copy Pair": false,
            "Is Writer": false
          }
        ]
      }
    },
    "Processing I/O": 0,
    "RAO Size": 0,
    "Minimum Processing Time": 0,
    "Sorting Type": "Section Reorder",
    "VVID": "00000001-04-00000003-01e46b88d6979fa4-3c73",
    "Mounted PV Name": "E0100100",
    "PV Name": "E0100100",
    "Mounted Device Id": 0,
    "Last Write": 1438876024,
    "Last Read": 1438876024,
    "Max Active Set": 1024,
    "Total Requests Scheduled": 0,
    "Active Set": 0,
    "Last Preference Change": 1438876024,
    "Write Preference Aid": 0,
    "Last Used": 95,
    "Number Scheduling Errors": 0,
    "Min Wait Time": 0,
    "Preference": "WRITERS",
    "I/O Efficiency": 0.0,
    "Average Wait Time": 0,
    "Max Wait Time": 0,
    "Average Schedule Time": 0,
    "Media Speed": 80,
    "Average Seek": 20000,
    "Number of Schedules": 0,
    "Schedule Time": 0,
    "Average Schedule": 10000,
  ]
}
```

```
"Number Unscheduled": 0,  
"Average Processing Time": 0,  
"Maximum Processing Time": 0  
}  
]  
}
```

17.2.2. JSON field descriptions

VV

The top level type is an ordered list, each containing statistics about a VV, its queue, and its sessions. The following member fields are defined:

Average Processing Time

The average time to process a request, in milliseconds.

VVID

The unique VV identifier

Average Seek

The average seek time, in milliseconds, for this media type. This is a static value set in the **Average Latency** field for the storage class.

Mounted PV Name

The cartridge label of the Mounted Volume. In a multi-stripe or RAIT VV, the "Mounted Volume" is the name of the volume which is being sent RAO requests.

Minimum Processing Time

The minimum time required to process a request, in milliseconds.

Number Scheduling Errors

The number of consecutive times a scheduling request has resulted in an error.

Average Schedule Time

The average time required to run the scheduler.

Schedule Time

Total time spent scheduling, in milliseconds.

Active Set

The currently active set number. Items which are not in the same set are never scheduled together; items in the lowest set are always scheduled first. The active set identifier is assigned to a request as requests enter the scheduler.

Max Active Set

The size of an active set. For example, when "1024" is specified up to 1024 requests will be scheduled together. Any items in addition to that 1024 will reside in a different active set, and will be scheduled after all items in the current set are complete.

PV Name

The first PV in the VV. This may not be the "Mounted PV" for multi-stripe VVs, but it will be how the VV is referenced in other places in the system.

Mounted Device Id

The device ID where the "Mounted PV" is mounted. If the PV is not mounted, this is zero.

Average Wait Time

Average time a request has been waiting before it is allowed to run, in milliseconds.

Sorting Type

The method used by the scheduler to order the requests. Valid values:

- `FIFO` : First in, First Out
- `Section Reorder` : Based upon the tape section and offset data
- `RAO` : Based upon an RAO (Recommended Access Order) request to the Mounted Device.

Last Used

Seconds since the last time the VV was used.

Last Write

Seconds since epoch of the last time a write was completed. Initially this will be the time when the queue was created.

Preference

The current queuing preference. Valid values:

- `WRITERS` indicates writers are preferred. This is the default value.
- `READERS` indicates readers are preferred. This preference may occur when readers have been waiting on a writer for a substantial amount of time.

I/O Efficiency

The number of megabytes per second transferred, compared to the media type's stated megabytes per second.

Note: This value is an estimate only and is biased low due to including mount and positioning times.

Maximum Processing Time

The maximum amount of time spent processing a request.

Last Read

Seconds since epoch of the last time a read was completed. Initially this will be the time when the queue was created.

Total Requests Scheduled

The total number of requests scheduled for this queue's lifetime. Note this is not a persistent value across mount requests, system outages, and so on.

Write Preference Aid

The activity identifier for the request which caused the queue to go into read preference mode. This allows the readers to run, but also allows the preference to switch to being write-preferenced after all long-running readers have completed.

Last Preference Change

Seconds since epoch of the last time the queuing preference changed.

Media Speed

The speed of the media, in megabytes per second. This is based upon the storage class.

Processing I/O

Total data processed divided by the total processing time.

RAO Size

Maximum RAO request size. For Section Ordered configurations which utilize RAO-capable devices, this represents the maximum number of requests that can be scheduled together (even if the Active Set is configured much higher). This is based upon the parameters of the media and is not configurable.

Min Wait Time

Minimum time a request has been waiting before it is allowed to run, in milliseconds.

Average Schedule

Average seek time, in milliseconds, when scheduling. This is currently statically defined as half of the storage class media's *average* latency.

Number of Schedules

Number of times the scheduler has been run.

Number Unscheduled

Number of requests waiting which have not yet been scheduled.

Queue

VV Queue Object containing details of requests.

Queue

Each VV has a Queue field, which contains details regarding the queue requests. The following member fields are defined:

Status

The queue's activity state. Valid values are:

- `Active` indicates the queue's session data is active.
- `Currently being rescored` indicates the queue's session data cannot be displayed since it is being rescored.

Sessions

An ordered list (ordered by Score) of all sessions in the queue and their details.

Sessions

Each request is associated with a session. The session data includes details about the request.

Start Section

The starting tape section for the request.

TimeOrder

The order of the request, in FIFO terms.

Score

The order of the request, based upon the Core Server (or VV's) ordering type.

Status

The status of the session. Valid values are:

- `RUNNING` indicates this request is currently running.
- `WAITING` indicates this request is currently waiting.

Activity Id

The unique activity ID for this session within the VV's queue.

Request Id

The request ID for this request.

Is Copy Pair

Whether this request belongs to a **copy pair** (usually a tape repack or tape-to-tape migration).

Wait Time

How long the request has been waiting, in milliseconds.

Data Length

Length of the data being read or written.

Start Offset

Starting section offset.

Is Complete

Whether the request is complete.

Is Reader

Whether the request is for a read.

Is Dismount

Whether the request is to dismount the tape.

Is Writer

Whether the request is a writer.

17.3. HPSS ACL permission exceptions

The purpose of this section is to describe the minor differences between the HPSS ACL implementation and the ACL implementation described in the *DFS Administration Guide*. The HPSS ACL implementation was modeled after the implementation described in the *DFS Administration Guide*, which can be found at

<http://www-01.ibm.com/software/stormgmt/dfs/library/docs/31manuals/AdminGd/duagd008.htm>

And with the exception of the d (delete) and i (insert) permissions, the HPSS ACL implementation is a duplication of the DFS implementation.

In the HPSS ACL implementation the d and i permissions can be set and deleted on directory objects (d and i permissions do not apply to files), but the existence (or absence) of d and i is ignored when determining whether or not a user has the permissions necessary to perform an operation.

For example HPSS allows objects to be created in a directory even if the i (insert) permission is not set on that directory. Likewise HPSS allows objects to be deleted from a directory even if that directory does not have the d (delete) permission set. In all other regards the HPSS ACL implementation mimics the DFS implementation.

Table 5 in the ACL Permission section of the *DFS Administration Guide* describes the various operations that can be performed on files or directories and the ACL permissions that are required to perform these operations. Table 5 is reproduced here with the d (delete) and i (insert) permissions removed. The following table describes how HPSS performs these operations and the ACL permissions required to perform them.

Operation	Required permissions
Change to a directory	Execute (x) on the directory itself. Execute (x) on all directories that lead to the directory.
List the contents of a directory	Read (r) on the directory itself. Execute (x) on all directories that lead to the directory.
List information about the objects in a directory	Read (r) and execute (x) on the directory itself. Execute (x) on all directories that lead to the directory.
Create an object	Write (w) , execute (x) on the directory in which the object is to be placed. Execute (x) on all directories that lead to the directory in which the object is to be placed.
Delete an object	Write (w) , execute (x) on the directory from which the object is to be deleted. Execute (x) on all directories that lead to the directory from which the object is to be deleted.
Rename an object	Write (w) , execute (x) on the object's current directory. Execute (x) on all directories that lead to the object's current directory. Write (w) , execute (x) on the object's new directory. Execute (x) on all directories that lead to the object's new directory. Write (w) on the object if the object is a directory.
Read or read lock a file	Read (r) on the file itself. Execute (x) on all directories that lead to the file.
Write or write lock a file	Write (w) on the file itself. Execute (x) on all directories that lead to the file.

Execute a binary file	Execute (x) on the file itself. Execute (x) on all directories that lead to the file.
Execute a shell script	Read (r) and execute (x) on the script itself. Execute (x) on all directories that lead to the script.
List the ACLs on an object	Execute (x) on all directories that lead to the object.
Change the ACLs on an object	Control (c) on the object itself. Execute (x) on all directories that lead to the object.

Appendix A. Glossary of terms and acronyms

ACL	Access Control List
ACSLs	Automated Cartridge System Library Software (Oracle StorageTek)
ADIC	Advanced Digital Information Corporation
accounting	The process of tracking system usage per user, possibly for the purposes of charging for that usage. Also, a log record type used to log accounting information.
ACI	AML Client Interface
AIX	Advanced Interactive Executive. An operating system provided on many IBM machines.
alarm	A log record type used to report situations that require administrator investigation or intervention.
AML	Automated Media Library. A tape robot.
AMS	Archive Management Unit
ANSI	American National Standards Institute
API	Application Program Interface
archive	One or more interconnected storage systems of the same architecture.
ASLR	Address Space Layout Randomization
attribute	When referring to a managed object, an attribute is one discrete piece of information, or set of related information, within that object.
attribute change	When referring to a managed object, an attribute change is the modification of an object attribute. This event may result in a notification being sent to SSM, if SSM is currently registered for that attribute.
audit (security)	An operation that produces lists of HPSS log messages whose record type is SECURITY. A security audit is used to provide a trail of security-relevant activity in HPSS.
AV	Account Validation
bar code	An array of rectangular bars and spaces in a predetermined pattern which represent alphanumeric information in a machine-readable format (such as a UPC symbol).
BFS	HPSS Bitfile Service
bitfile	A file stored in HPSS, represented as a logical string of bits unrestricted in size or internal structure. HPSS imposes a size limitation in 8-bit bytes based upon the maximum size in bytes that can be represented by a 64-bit unsigned integer.
bitfile segment	An internal metadata structure, not normally visible, used by the Core Server to map contiguous pieces of a bitfile to underlying storage.
Bitfile Service	Portion of the HPSS Core Server that provides a logical abstraction of bitfiles to its clients.

BBTM	Blocks Between Tape Marks. The number of data blocks that are written to a tape virtual volume before a tape mark is required on the physical media.
CAP	Cartridge Access Port
cartridge	A physical media container, such as a tape reel or cassette, capable of being mounted on and dismounted from a drive. A fixed disk is technically considered to be a cartridge because it meets this definition and can be logically mounted and dismounted.
class	A type definition in Java. It defines a template on which objects with similar characteristics can be built, and includes variables and methods specific to the class.
Class of Service	A set of storage system characteristics used to group bitfiles with similar logical characteristics and performance requirements together. A Class of Service is supported by an underlying hierarchy of storage classes.
cluster	The unit of storage space allocation on HPSS disks. The smallest amount of disk space that can be allocated from a virtual volume is a cluster. The size of the cluster on any given disk volume is determined by the size of the smallest storage segment that will be allocated on the volume, and other factors.
configuration	The process of initializing or modifying various parameters affecting the behavior of an HPSS server or infrastructure service.
COS	Class of Service
control path	For the SCSI PVR, this is a connection to the library for sending commands. Control paths can be discovered using <code>device_scan</code> .
Core Server	An HPSS server which manages the name space and storage for an HPSS system. The Core Server manages the name space in which files are defined, the attributes of the files, and the storage media on which the files are stored. The Core Server is the central server of an HPSS system. Each storage subsystem uses exactly one Core Server.
CRC	Cyclic Redundancy Check
CS	Core Server
daemon	A UNIX program that runs continuously in the background.
DAS	Distributed AML Server
DB2	A relational database system, a product of IBM Corporation, used by HPSS to store and manage HPSS system metadata.
DCE	Distributed Computing Environment
debug	A log record type used to report internal events that can be helpful in troubleshooting the system.
delog	The process of extracting, formatting, and outputting HPSS central log records. This process is obsolete in 7.4 and later versions of HPSS. HPSS logs are now recorded as plain text.
deregistration	The process of disabling notification to SSM for a particular attribute change.
descriptive name	A human-readable name for an HPSS server.
device	A physical piece of hardware, usually associated with a drive, that is capable of reading or writing data.

directory	An HPSS object that can contain files, symbolic links, hard links, and other directories.
dismount	An operation in which a cartridge is either physically or logically removed from a device, rendering it unreadable and unwritable. In the case of tape cartridges, a dismount operation is a physical operation. In the case of a fixed disk unit, a dismount is a logical operation.
DNS	Domain Name Service
DOE	Department of Energy
DPF	Database Partitioning Feature
drive	A physical piece of hardware capable of reading or writing mounted cartridges. The terms device and drive are often used interchangeably.
EB	Exabyte (2^{60})
EOF	End of File
EOM	End of Media
ERA	Extended Registry Attribute
ESCON	Enterprise System Connection
event	A log record type used to report informational messages (for example, subsystem starting or subsystem terminating).
export	An operation in which a cartridge and its associated storage space are removed from the HPSS system Physical Volume Library. It may or may not include an eject, which is the removal of the cartridge from its Physical Volume Repository.
FC SAN	Fiber Channel Storage Area Network
FIFO	First in first out
file	An object than can be written to, read from, or both, with attributes including access permissions and type, as defined by POSIX (P1003.1-1990). HPSS supports only regular files.
file family	An attribute of an HPSS file that is used to group a set of files on a common set of tape virtual volumes.
fileset	A collection of related files that are organized into a single easily managed unit. A fileset is a disjoint directory tree that can be mounted in some other directory tree to make it accessible to users.
fileset ID	A 64-bit number that uniquely identifies a fileset.
fileset name	A name that uniquely identifies a fileset.
file system ID	A 32-bit number that uniquely identifies an aggregate.
FTP	File Transfer Protocol
FSF	Forward Space File
FSR	Forward Space Record
Gatekeeper	An HPSS server that provides two main services: the ability to schedule the use of HPSS resources referred to as the Gatekeeping Service, and the ability to validate user accounts referred to as the Account Validation Service.

Gatekeeping Service	A registered interface in the Gatekeeper that provides a site the mechanism to create local policy on how to throttle or deny create, open and stage requests and which of these request types to monitor.
Gatekeeping Site Interface	The APIs of the gatekeeping site policy code.
Gatekeeping Site Policy	The gatekeeping shared library code written by the site to monitor and throttle create, open, and/or stage requests.
GB	Gigabyte (2^{30})
GECOS	The comment field in a UNIX password entry that can contain general information about a user, such as office or phone number.
GID	Group Identifier
GK	Gatekeeper
GSS	Generic Security Service
GUI	Graphical User Interface
HA	High Availability
HACMP	High Availability Clustered Multi-Processing - A software package used to implement high availability systems.
HADR	DB2 High Availability Disaster Recovery
halt	A forced shutdown of an HPSS server.
HBA	Host Bus Adapter
HDM	Shorthand for HPSS/DMAP.
hierarchy	See <i>storage hierarchy</i> .
HPSS	High Performance Storage System
HPSS-only fileset	An HPSS fileset that is not linked to an external file system (such as XFS).
HTP	HPSS Test Plan
IBM	International Business Machines Corporation
ID	Identifier
IDE	Integrated Drive Electronics
IEEE	Institute of Electrical and Electronics Engineers
import	An operation in which a cartridge and its associated storage space are made available to the HPSS system. An import requires that the cartridge has been physically introduced into a Physical Volume Repository (injected). Importing the cartridge makes it known to the Physical Volume Library.
I/O	Input/Output
IOD/IOR	I/O Descriptor/I/O Reply. Structures used to send control information about data movement requests in HPSS and about the success or failure of the requests.
IP	Internet Protocol
IRIX	SGI's implementation of UNIX
JRE	Java Runtime Environment
junction	A mount point for an HPSS fileset.

KB	Kilobyte (2^{10})
KDC	Key Distribution Center
LAN	Local Area Network
LANL	Los Alamos National Laboratory
latency	For tape media, the average time in seconds between the start of a read or write request and the time when the drive actually begins reading or writing the tape.
LBP	Logical Block Protection
LDAP	Lightweight Directory Access Protocol
LFT	Local File Transfer
LLNL	Lawrence Livermore National Laboratory
LMU	Library Management Unit
Location Server	An HPSS server that is used to help clients locate the appropriate Core Server or other HPSS server to use for a particular request.
log record	A message generated by an HPSS application and handled and recorded by the HPSS logging subsystem.
log record type	A log record may be of type alarm, event, info, debug, request, security, trace, or accounting.
logging service	An HPSS infrastructure service consisting of the logging subsystem and one or more logging policies. A default logging policy can be specified, which will apply to all servers, or server-specific logging policies may be defined.
LS	Location Server
LSM	Library Storage Module
LTO	Linear Tape-Open. A half-inch open tape technology developed by IBM, HP, and Seagate.
LUN	Logical Unit Number
LVM	Logical Volume Manager
MAC	Mandatory Access Control
managed object	A programming data structure that represents an HPSS system resource. The resource can be monitored and controlled by operations on the managed object. Managed objects in HPSS are used to represent servers, drives, storage media, jobs, and other resources.
MB	Megabyte (2^{20})
MBS	Media Block Size
metadata	Control information about the data stored under HPSS, such as location, access times, permissions, and storage policies. Most HPSS metadata is stored in a DB2 relational database.
method	A Java function or subroutine.
migrate	To copy file data from a level in the file's hierarchy onto the next lower level in the hierarchy.
Migration/Purge Server	An HPSS server responsible for supervising the placement of data in the storage hierarchies based upon site-defined migration and purge policies.

MM	Metadata Manager. A software library that provides a programming API to interface HPSS servers with the DB2 programming environment.
mount	An operation in which a cartridge is either physically or logically made readable/writable on a drive. In the case of tape cartridges, a mount operation is a physical operation. In the case of a fixed disk unit, a mount is a logical operation.
mount point	A place where a fileset is mounted in the XFS and HPSS name spaces.
Mover	An HPSS server that provides control of storage devices and data transfers within HPSS.
MPS	Migration/Purge Server
MVR	Mover
NASA	National Aeronautics and Space Administration
Name Service	The portion of the Core Server that provides a mapping between names and machine-oriented identifiers. In addition, the Name Service performs access verification and provides the Portable Operating System Interface (POSIX).
name space	The set of name-object pairs managed by the HPSS Core Server.
NERSC	National Energy Research Supercomputer Center
NIS	Network Information Service
NLS	National Language Support
notification	A notice from one server to another about a noteworthy occurrence. HPSS notifications include notices sent from other servers to SSM of changes in managed object attributes, changes in tape mount information, and log messages of type alarm or event.
NS	HPSS Name Service
NSL	National Storage Laboratory
object	See <i>managed object</i> .
ORNL	Oak Ridge National Laboratory
OS	Operating System
OS/2	The operating system (multi-tasking, single user) used on the AMU controller PC.
PB	Petabyte (2^{50})
PFTP	Parallel File Transfer Protocol
PFTPD	PFTP Daemon
physical volume	An HPSS object managed jointly by the Core Server and the Physical Volume Library that represents the portion of a virtual volume. A virtual volume may be composed of one or more physical volumes, but a physical volume may contain data from no more than one virtual volume.
Physical Volume Library	An HPSS server that manages mounts and dismounts of HPSS physical volumes.
Physical Volume Repository	An HPSS server that manages the robotic agent responsible for mounting and dismounting cartridges or interfaces with the human agent responsible for mounting and dismounting cartridges.
PIO	Parallel I/O

PIOFS	Parallel I/O File System
POSIX	Portable Operating System Interface (for computer environments).
purge	Deletion of file data from a level in the file's hierarchy after the data has been duplicated at lower levels in the hierarchy and is no longer needed at the deletion level.
purge lock	A lock applied to a bitfile which prohibits the bitfile from being purged.
PV	Physical Volume
PVL	Physical Volume Library
PVM	Physical Volume Manager
PVR	Physical Volume Repository
RAID	Redundant Array of Independent Disks
RAIT	Redundant Array of Independent Tapes
RAM	Random Access Memory
RAO	Recommended Access Order
reclaim	The act of making previously written but now empty tape virtual volumes available for reuse. Reclaimed tape virtual volumes are assigned a new Virtual Volume ID, but retain the rest of their previous characteristics. Reclaim is also the name of the utility program that performs this task.
registration	The process by which SSM requests notification of changes to specified attributes of a managed object.
reinitialization	An HPSS SSM administrative operation that directs an HPSS server to reread its latest configuration information, and to change its operating parameters to match that configuration, without going through a server shutdown and restart.
repack	The act of moving data from a virtual volume onto another virtual volume with the same characteristics with the intention of removing all data from the source virtual volume. Repack is also the name of the utility program that performs this task.
request	A log record type used to report some action being performed by an HPSS server on behalf of a client.
RISC	Reduced Instruction Set Computer/Cycles
RPC	Remote Procedure Call
RSF	Reverse Space File
RSR	Reverse Space Record
SCSI	Small Computer Systems Interface
security	A log record type used to report security-related events (for example, authorization failures).
SGI	Silicon Graphics
shelf tape	A cartridge which has been physically removed from a tape library but whose file metadata still resides in HPSS.
shutdown	An HPSS SSM administrative operation that causes a server to stop its execution gracefully.
sink	The set of destinations to which data is sent during a data transfer, such as disk devices, memory buffers, or network addresses.

SM	System Manager
SMC	SCSI Medium Changer
SME	Subject Matter Expert
SNL	Sandia National Laboratories
SOID	Storage Object ID. An internal HPSS storage object identifier that uniquely identifies a storage resource. The SOID contains a unique identifier for the object, and a unique identifier for the server that manages the object.
source	The set of origins from which data is received during a data transfer, such as disk devices, memory buffers, or network addresses.
SP	Scalable Processor
SS	HPSS Storage Service
SSD	Solid State Drive
SSH	Secure Shell
SSI	Storage Server Interface
SSM	Storage System Management
SSM session	The environment in which an SSM user interacts with the SSM System Manager to monitor and control HPSS. This environment may be the graphical user interface provided by the hpssgui program, or it may be the command-line user interface provided by the hpssadm program.
SSMSM	Storage System Management System Manager
stage	To copy file data from a level in the file's hierarchy onto the top level in the hierarchy.
start-up	An HPSS SSM administrative operation that causes a server to begin execution.
info	A log record type used to report file staging and other kinds of information.
STK	Storage Technology Corporation (Oracle StorageTek)
storage class	An HPSS object used to group storage media together to provide storage for HPSS data with specific characteristics. The characteristics are both physical and logical.
storage hierarchy	An ordered collection of storage classes. The hierarchy consists of a fixed number of storage levels numbered from level 1 to the number of levels in the hierarchy, with the maximum level being limited to 5 by HPSS. Each level is associated with a specific storage class. Migration and stage commands result in data being copied between different storage levels in the hierarchy. Each Class of Service has an associated hierarchy.
storage level	The relative position of a single storage class in a storage hierarchy. For example, if a storage class is at the top of a hierarchy, the storage level is 1.
storage map	An HPSS object managed by the Core Server to keep track of allocated storage space.
storage segment	An HPSS object managed by the Core Server to provide abstract storage for a bitfile or parts of a bitfile.
Storage Service	The portion of the Core Server which provides control over a hierarchy of virtual and physical storage resources.
storage subsystem	A portion of the HPSS name space that is managed by an independent Core Server and (optionally) Migration/Purge Server.

Storage System Management	An HPSS component that provides monitoring and control of HPSS via a windowed operator interface or command-line interface.
stripe length	The number of bytes that must be written to span all the physical storage media (physical volumes) that are grouped together to form the logical storage media (virtual volume). The stripe length equals the virtual volume block size multiplied by the number of physical volumes in the stripe group (that is, stripe width).
stripe width	The number of physical volumes grouped together to represent a virtual volume.
System Manager	The Storage System Management (SSM) server. It communicates with all other HPSS components requiring monitoring or control. It also communicates with the SSM graphical user interface (hpssgui) and command line interface (hpssadm).
TB	Terabyte (2^{40})
TCP/IP	Transmission Control Protocol/Internet Protocol
TDS	Tivoli Directory Server
TI-RPC	Transport-Independent-Remote Procedure Call
trace	A log record type used to record procedure entry/exit events during HPSS server software operation.
transaction	<p>A programming construct that enables multiple data operations to possess the following properties:</p> <ul style="list-style-type: none">• All operations commit or abort/roll-back together such that they form a single unit of work.• All data modified as part of the same transaction are guaranteed to maintain a consistent state whether the transaction is aborted or committed.• Data modified from one transaction are isolated from other transactions until the transaction is either committed or aborted.• Once the transaction commits, all changes to data are guaranteed to be permanent.
TSA/MP	Tivoli System Automation for Multiplatforms
TSM	Tivoli Storage Manager
UDA	User-defined Attribute
UDP	User Datagram Protocol
UID	User Identifier
UPC	Universal Product Code
UUID	Universal Unique Identifier
VPN	Virtual Private Network
virtual volume	<p>An HPSS object managed by the Core Server that is used to represent logical media. A virtual volume is made up of a group of physical storage media (a stripe group of physical volumes).</p>
virtual volume block size	The size of the block of data bytes that is written to each physical volume of a striped virtual volume before switching to the next physical volume.

VV	Virtual Volume
XDSM	The Open Group's Data Storage Management standard. It defines APIs that use events to notify Data Management applications about operations on files.
XFS	A file system created by SGI available as open source for the Linux operating system.
XML	Extensible Markup Language

Appendix B. References

1. **3580 Ultrium Tape Drive Setup, Operator and Service Guide** GA32-0415-00
2. **3584 UltraScalable Tape Library Planning and Operator Guide** GA32-0408-01
3. **3584 UltraScalable Tape Library SCSI Reference** WB1108-00
4. **AIX Performance Tuning Guide**
5. **Data Storage Management (XDSM) API**, ISBN 1-85912-190-X
6. **HACMP for AIX, Version 4.4: Concepts and Facilities**
7. **HACMP for AIX, Version 4.4: Planning Guide**
8. **HACMP for AIX, Version 4.4: Installation Guide**
9. **HACMP for AIX, Version 4.4: Administration Guide**
10. **HACMP for AIX, Version 4.4: Troubleshooting Guide**
11. **HPSS Error Messages Reference Manual**, current release
12. **HPSS Programmer's Reference**, current release
13. **HPSS Programmer's Reference - I/O Supplement**, current release
14. **HPSS User's Guide**, current release
15. **IBM SCSI Device Drivers: Installation and User's Guide**, GC35-0154-01
16. **IBM Ultrium Device Drivers Installation and User's Guide** GA32-0430-00.1
17. **IBM Ultrium Device Drivers Programming Reference** WB1304-01
18. **Interfacing Guide DAS**, Order no. DOC F00 011
19. **Installing, Managing, and Using the IBM AIX Parallel I/O File System**, SH34-6065-02
20. **Platform Notes: The hme FastEthernet Device Driver** 805-4449
21. **POSIX 1003.1-1990 Tar Standard**
22. **Reference Guide AMU**, Order no. DOC E00 005
23. **STK Automated Cartridge System Library Software (ACSLs) System Administrator's Guide**, PN 16716
24. **STK Automated Cartridge System Library Software Programmer's Guide**, PN 16718

- 25.J. Steiner, C. Neuman, and J. Schiller, "**Kerberos: An Authentication Service for Open Network Systems**," USENIX 1988 Winter Conference Proceedings (1988).
- 26.R.W. Watson and R.A. Coyne, "**The Parallel I/O Architecture of the High-Performance Storage System (HPSS)**," from the 1995 IEEE MSS Symposium, courtesy of the IEEE Computer Society Press.
- 27.T.W. Tyler and D.S. Fisher, "**Using Distributed OLTP Technology in a High-Performance Storage System**," from the 1995 IEEE MSS Symposium, courtesy of the IEEE Computer Society Press.
- 28.J.K. Deutsch and M.R. Gary, "**Physical Volume Library Deadlock Avoidance in a Striped Media Environment**," from the 1995 IEEE MSS Symposium, courtesy of the IEEE Computer Society Press.
- 29.R. Grossman, X. Qin, W. Xu, H. Hulen, and T. Tyler, "**An Architecture for a Scalable, High-Performance Digital Library**," from the 1995 IEEE MSS Symposium, courtesy of the IEEE Computer Society Press.
- 30.S. Louis and R.D. Burris, "**Management Issues for High-Performance Storage Systems**," from the 1995 IEEE MSS Symposium, courtesy of the IEEE Computer Society Press.
- 31.D. Fisher, J. Sobolewski, and T. Tyler, "**Distributed Metadata Management in the High Performance Storage System**," from the 1st IEEE Metadata Conference, April 16-18, 1996.

Appendix C. Developer acknowledgments

HPSS is a product of a government-industry collaboration. The project approach is based on the premise that no single company, government laboratory, or research organization has the ability to confront all of the system-level issues that must be resolved for significant advancement in high-performance storage system technology.

HPSS development was performed jointly by IBM Worldwide Government Industry, Lawrence Berkeley National Laboratory, Lawrence Livermore National Laboratory, Los Alamos National Laboratory, NASA Langley Research Center, Oak Ridge National Laboratory, and Sandia National Laboratories.

We would like to acknowledge Argonne National Laboratory, the National Center for Atmospheric Research, and Pacific Northwest Laboratory for their help with initial requirements reviews.

We also wish to acknowledge Cornell Information Technologies of Cornell University for providing assistance with naming service and transaction management evaluations and for joint developments of the Name Service.

In addition, we wish to acknowledge the many discussions, design ideas, implementation and operation experiences we have shared with colleagues at the National Storage Laboratory, the IEEE Mass Storage Systems and Technology Technical Committee, the IEEE Storage System Standards Working Group, and the storage community at large.

We also wish to acknowledge the Cornell Theory Center and the Maui High Performance Computer Center for providing a test bed for the initial HPSS release.

We also wish to acknowledge Gleicher Enterprises, LLC for the development of the HSI, HTAR, and Transfer Agent client applications.

Finally, we wish to acknowledge CEA-DAM (**Commissariat à l'Énergie Atomique - Centre d'Études de Bruyères-le-Châtel**) for providing assistance with development of NFS V3 protocol support.