HITACHI
Inspire the Next

Hitachi Command Suite

User Guide

Hitachi Data Systems

MK-90HC172-25

© 2014, 2016 Hitachi, Ltd. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including copying and recording, or stored in a database or retrieval system for commercial purposes without the express written permission of Hitachi, Ltd., or Hitachi Data Systems Corporation (collectively "Hitachi"). Licensee may make copies of the Materials provided that any such copy is: (i) created as an essential step in utilization of the Software as licensed and is used in no other manner; or (ii) used for archival purposes. Licensee may not make any other copies of the Materials. "Materials" mean text, data, photographs, graphics, audio, video and documents.

Hitachi reserves the right to make changes to this Material at any time without notice and assumes no responsibility for its use. The Materials contain the most current information available at the time of publication.

Some of the features described in the Materials might not be currently available. Refer to the most recent product announcement for information about feature and product availability, or contact Hitachi Data Systems Corporation at https://support.hds.com/en_us/contact-us.html.

Notice: Hitachi products and services can be ordered only under the terms and conditions of the applicable Hitachi agreements. The use of Hitachi products is governed by the terms of your agreements with Hitachi Data Systems Corporation.

By using this software, you agree that you are responsible for:

1. Acquiring the relevant consents as may be required under local privacy laws or otherwise from authorized employees and other individuals to access relevant data; and
2. Verifying that data continues to be held, retrieved, deleted, or otherwise processed in accordance with relevant laws.

Notice on Export Controls. The technical data and technology inherent in this Document may be subject to U.S. export control laws, including the U.S. Export Administration Act and its associated regulations, and may be subject to export or import regulations in other countries. Reader agrees to comply strictly with all such regulations and acknowledges that Reader has the responsibility to obtain licenses to export, re-export, or import the Document and any Compliant Products.

Hitachi is a registered trademark of Hitachi, Ltd., in the United States and other countries.

AIX, AS/400e, DB2, Domino, DS6000, DS8000, Enterprise Storage Server, eServer, FICON, FlashCopy, IBM, Lotus, MVS, OS/390, PowerPC, RS/6000, S/390, System z9, System z10, Tivoli, z/OS, z9, z10, z13, z/VM, and z/VSE are registered trademarks or trademarks of International Business Machines Corporation.

Active Directory, ActiveX, Bing, Excel, Hyper-V, Internet Explorer, the Internet Explorer logo, Microsoft, the Microsoft Corporate Logo, MS-DOS, Outlook, PowerPoint, SharePoint, Silverlight, SmartScreen, SQL Server, Visual Basic, Visual C++, Visual Studio, Windows, the Windows logo, Windows Azure, Windows PowerShell, Windows Server, the Windows start button, and Windows Vista are registered trademarks or trademarks of Microsoft Corporation. Microsoft product screen shots are reprinted with permission from Microsoft Corporation.

All other trademarks, service marks, and company names in this document or website are properties of their respective owners.

Contents

Preface..... 15
Intended audience..... 16
Product version..... 16
Release notes..... 16
Related documents..... 16
Document conventions..... 16
Conventions for storage capacity values..... 17
Accessing product documentation..... 18
Getting help..... 18
Comments..... 19

1 Overview of Hitachi Command Suite..... 21
About Hitachi Command Suite..... 22
Features..... 23
What's new..... 24
System configuration..... 25
Process flow..... 28
Navigating the interface..... 29
Navigating help..... 31

2 Setting up Hitachi Command Suite..... 33
Configuring Hitachi Command Suite..... 34
Configuring your browser and Java for Hitachi Command Suite..... 34
About configuring browser settings..... 34
Checking management server name resolution..... 35
Disabling pop-up blocking..... 36
Setting options for Internet Explorer..... 37
Setting options for Firefox..... 38
Setting options for Google Chrome..... 39
Specifying settings to link with Device Manager - Storage Navigator..... 39
Setting the Java™ Web Start (versions 6.0 and 7.0) proxy to link with other products..... 39
Configuring JRE versions from JWS version 6.0..... 40
Clearing the cache when upgrading Hitachi Command Suite..... 41

Logging in to Hitachi Command Suite..... 41
Logging in when HCS is not available.....42
Setting up security..... 43
About configuring security options..... 44
Setting a password policy.....44
Setting automatic account locking.....45
Setting a warning banner message..... 45
Downloading components.....46
About downloading components..... 46
Downloading agents, CLI, and Host Data Collector files..... 46
Managing HCS licenses.....46
About HCS license management..... 46
Registering an HCS license..... 47
Checking HCS license information..... 48
Managing licenses for Virtual Storage Platform storage systems.....48
License key types..... 48
Installing the software..... 49
Enabling a license.....50
Disabling a license..... 50
Removing a software application.....51
Updating license status..... 51
Viewing license information..... 52
Examples of license information.....52

3 Discovering, registering, and adding management targets.....55

Setting up storage resources.....56
Registering storage systems.....56
About registering and removing a storage system..... 57
Prerequisites for registering a storage system..... 57
Registering a storage system.....60
Changing storage system information.....62
Removing a storage system.....63
About acquiring the most recent storage system information..... 63
Acquiring the most recent storage system information.....64
Operations available to SMI-S enabled storage systems..... 65
Registering hosts..... 66
About registering a host.....66
Methods for registering hosts..... 66
Priority for acquiring the WWN or iSCSI name.....68
Registering hosts by using Host Data Collector.....69
Registering hosts manually by specifying the WWN or iSCSI name..... 70
Registering hosts using host scan..... 71
Registering hosts using merge hosts..... 72
About changing host settings and information.....73
Workflow for detecting hosts..... 74
Updating host information registered by using Host Data Collector..... 75
Changing settings for a manually registered host..... 76
Changing settings for a host registered by using Device Manager agent.....77
Refreshing vCenter hosts.....78
About removing hosts and related resources, or vCenters.....79
Removing hosts and releasing associated resources..... 80

Registering file servers..... 81
About registering and removing file servers..... 81
Environment settings for NAS Platform family..... 82
Registering file servers..... 84
Changing the name of a file server..... 85
Changing the Admin services EVS settings of NAS Platform..... 86

4 Setting up users and access control..... 87

Setting up users and access control..... 88
About user accounts and controlling access to resources..... 88
Creating and managing user accounts..... 91
Creating a user account..... 91
User ID and password policies..... 92
Editing the profile for a user account..... 93
Editing your own user profile..... 94
Changing the password for a user account..... 94
Changing your own password..... 94
Changing permissions for a user account..... 95
Changing the lock status of user accounts..... 95
Configuring external authentication for users..... 96
Configuring external authentication for groups..... 97
Deleting user accounts..... 98
Controlling access to resources..... 99
About access control..... 99
Access control examples..... 100
About resource groups..... 102
Prerequisites for creating resource groups..... 104
Creating resource groups..... 104
Editing a resource group..... 105
Deleting resource groups..... 106
About user groups..... 107
User group roles..... 110
Custom roles..... 111
Required roles and resource groups by function..... 114
Creating user groups..... 134
Editing a user group..... 135
Assigning resource groups and roles to a user group..... 136
Changing a user's user group..... 136
Deleting user groups..... 136

5 Provisioning storage..... 139

Creating a storage operating environment..... 140
Allocating storage..... 141
Creating and deleting parity groups..... 142
About creating parity groups..... 142
Creating parity groups..... 143
Create parity groups dialog box..... 144
About deleting parity groups..... 145
Deleting parity groups..... 145
Spare drives and copy-back mode..... 146

Assigning spare drives....147
Releasing spare drives....148
Editing copy-back mode....148
Enabling accelerated compression....149
Creating and deleting volumes....149
About creating volumes....149
Notes on performing quick formats....151
Creating volumes....151
Create Volumes dialog box....152
About shredding volume data....155
Shredding volume data....155
About deleting unallocated volumes....156
Deleting unallocated volumes....157
About creating a LUSE volume....157
Creating a LUSE volume....158
About releasing a LUSE volume....159
Releasing a LUSE volume....159
Virtualizing external storage....160
About virtualizing and unvirtualizing volumes....160
Virtualizing volumes of a registered storage system....162
Discovering and virtualizing volumes of an unregistered storage system....163
Unvirtualizing volumes....165
Virtualize Volumes dialog box....166
Virtualizing storage capacity (HDP/HDT)....171
About virtualizing storage capacity....171
Creating a DP pool....173
Create Pool dialog box....176
Verifying DP pool information....180
Expanding DP pools....180
Shrinking a DP pool....183
Modifying DP pool settings....184
Deleting DP pools....184
Expanding DP volumes....185
Reclaiming zero pages....186
Changing an access attribute to read/write....186
Virtualizing storage tiers (HDT)....187
About virtualizing storage tiers....187
Manually starting or stopping the monitoring of HDT pools....190
Manually starting or stopping the tier relocation of an HDT pool....191
Scheduling monitoring and tier relocation of HDT pools....191
Editing tier relocation for HDT volumes....192
Applying a tiering policy to HDT volumes....193
Customizing a tiering policy for HDT volumes....193
Notes on data placement profiles for HDT volumes....194
Creating a data placement profile for HDT volumes....195
Updating a data placement profile for HDT volumes....196
Editing a data placement profile for HDT volumes....197
Applying a data placement profile for HDT volumes....198
Scheduling data placement profiles for HDT volumes....199
Editing an external LDEV tiering rank for an HDT pool....200
Allocating and unallocating volumes....200
About allocating volumes....201

Volume allocation methods....202
Prerequisites for allocating volumes....204
Allocating volumes from general tasks....204
Allocating volumes to selected hosts....205
Allocating volumes to selected file servers....206
Allocating selected volumes to hosts....208
Allocating volumes to clustered hosts....209
Allocating volumes by using a keyword search....210
Allocating volumes by using a criteria search....211
Allocating volumes by using existing volume settings....212
Allocate Volumes dialog box....213
About clustered-host storage....223
Creating clustered-host storage....224
About unallocating volumes....226
Unallocating volumes from hosts....227
Unallocating volumes from file servers....228
Unallocate volumes dialog box....229
Configuring Fibre Channel and iSCSI ports....231
Enabling LUN security on a port....231
Disabling LUN security on a port....232
Editing iSCSI port settings....233
Setting the T10 PI mode on a port....234
Setting the data transfer speed on a Fibre Channel port....235
Setting the Fibre Channel port address....237
Setting the fabric switch....237
Managing LUN Paths....238
About LUN path management....238
Editing LUN paths....240
Editing the host mode and host mode options....242
Editing LUN paths when exchanging a failed HBA....243
Editing LUN paths when adding or exchanging an HBA....244
Removing LUN paths after adding an HBA....245
Managing NAS Platform file systems and shares....246
About managing NAS Platform file systems and shares....246
Workflow for providing NAS Platform file shares....250
Conditions for parity groups that are used in storage pools....253
Conditions for DP pools that are used in storage pools....253
Best practices for configuring storage pools....254
Notes on NAS Platform file system capacity....256
Notes on setting LUN paths for NAS Platform....257
Creating a NAS Platform storage pool (v11.3 or later)....257
Expanding a NAS Platform storage pool (v11.3 or later)....259
Changing a NAS Platform module storage pool name....260
Creating a NAS Platform storage pool (earlier than v11.3)....261
Expanding a NAS Platform storage pool (earlier than v11.3)....262
Changing NAS Platform storage pool settings....263
Deleting NAS Platform storage pools....264
Creating a NAS Platform file system....264
Expanding a NAS Platform file system....266
Changing NAS Platform file system settings....267
Mounting NAS Platform file systems....268
Unmounting NAS Platform file systems....268

Deleting a NAS Platform file system....269
Adding NAS Platform file shares....270
Changing NAS Platform file share settings....271
Deleting NAS Platform file shares....272
Managing NAS Platform F and Data Ingestor file systems and shares....273
About managing NAS Platform F and Data Ingestor file systems and shares....273
Workflow for providing HNAS F and HDI file shares (HFSM)....275
Creating HNAS F or HDI file systems (HFSM v3.2 or later)....276
Creating HDI file systems (HFSM v3.1)....277
Changing HNAS F or HDI file system settings (HFSM v3.1 or later)....278
Deleting HNAS F or HDI file systems....279
File share configuration examples....279
Managing copy pairs....281
About pair management....281
HCS copy pair management operations (no Replication Manager license)....281
Defining copy pairs....282
Changing the status of a copy pair....283
Managing command devices....283
Adding command devices....283
Editing command devices....284
Deleting command devices....285

6 Managing replication configurations using the Replication tab....287

About replicating volumes....288
HCS Replication tab operations (with Replication Manager license)....288
Managing replication configurations....292
Volume replication configuration workflow....292
Creating a replication environment....293
Creating remote paths for a replication configuration....294
Creating DP pools for a replication configuration....295
Configuring journal groups for a replication configuration....296
Creating Thin Image pools for a replication configuration....297
Configuring pair management servers for a replication configuration....297
Configuring a quorum disk for a replication configuration....298
Configuring a virtual storage machine for a replication configuration....299
Discontinuing a replication configuration....300
Site management with the Replication tab....302
Managing copy topologies....303
Creating a copy topology for a replication configuration (TC/UR)....303
Creating a copy topology for a replication configuration (GAD)....305
Changing the copy topology name....306
About changing the copy topology type....307
Expanding a copy topology by changing the type....308
Shrinking a copy topology by changing the type....309
Expanding a copy topology by adding a copy pair....310
Expanding a copy topology by adding a local copy group....311
Expanding an unsupported copy topology....312
Shrinking a copy topology by deleting a copy pair....312
Shrinking a copy topology by deleting a local copy group....313
Adding and deleting copy pairs in an unsupported copy topology....313
Deleting a copy topology from a replication configuration....314

Monitoring replication configurations....314
Setting alerts for a replication configuration....314
Monitoring the status of a replication configuration....315
Changing the status of a replication configuration....316
Changing the status of an unsupported configuration....317
Replication and disaster recovery....317
Backing up and restoring a local copy....317
Performing replication takeover and takeback....318
Performing a disaster recovery test using a backup (split TC/UR copy)....318
Performing a disaster recovery test using a snapshot (resync and split SI/TI copy) 319

7 Managing virtual storage machines....321

About virtual storage machines....322
Prerequisites for allocating resources to virtual storage machines....326
Creating virtual storage machines....327
Allocating volumes to hosts by using virtual storage machine resources....328
Editing virtual storage machines....330
Deleting virtual storage machines....331

8 Replicating volumes using global-active device 333

About global-active device....334
Setting up a global-active device environment....340
About global-active device setup....340
Prerequisites for global-active device setup....342
Select primary and secondary storage systems....342
Configuring remote paths....343
Configuring the quorum disk....345
Configuring pair management servers....348
Configuring the virtual storage machine....349
Monitoring and managing global-active device pairs....350
Allocating global-active device pairs....350
Allocating global-active device pairs based on like volumes....353
Allocating global-active device pairs using allocated volumes....356
Checking global-active device pair status....358
Monitoring global-active device pair performance....359
Unallocating global-active device pairs....360
Unallocating individual global-active device volumes....362
Recovering from global-active device failures....363
About global-active device failures....363
Recovering from HBA-related host path failures....364
Discontinuing a global-active device environment....365

9 Migrating data using nondisruptive migration (VSP G1000, VSP G1500, VSP F1500 and VSP Gx00 models)....367

About nondisruptive migration....368
About migration planning....369
Setting up migration projects and jobs....370
Understanding migration projects and jobs....370

Navigating migration projects and jobs....371
Preparing storage resources for migration....375
Creating a migration project....376
Creating a migration job....378
Migrating volumes to the target storage system....380
Volume and copy pair configuration migration workflow....380
Step 1: Virtualize and Allocate Volumes....382
Step 2: Change host I/O paths....383
Step 3: Unallocate source volumes....385
Step 4: Nondisruptive migration....386
Step 5: Unvirtualize source volumes....387
Migrating copy pair configurations to the target storage system....388
Preparing target secondary volumes....388
Allocating target secondary volumes....390
Migrating configuration definition files....390
Creating copy pairs on the target storage system....391
Managing migration projects and jobs....393
Tracking migration project status....393
Tracking migration job status....394
About resuming failed migration jobs....396
Changing the status of a failed task....396
Editing a migration project....398
Editing a migration job....399
Canceling a migration job....399
Canceling a volume from a migration job....400
Deleting a migration project....400
Post-migration tasks....401
Shredding or deleting source volumes....401
Removing the source storage system from Hitachi Command Suite....402
Reclaiming zero pages for target volumes....402
Optimizing the target storage configuration....403

10 Configuring a VVol environment....405

Working with VMware VVols....406
Storage Policy Based Management (SPBM)....407
About Storage Policy Based Management....407
About default profiles....408
SPBM implementation examples....408
VVols and block storage....409
VVol configuration and operation tasks for block storage....409
Reviewing supported functions for VVol management....413
Guidelines for working with virtual machines in a VMware VVol environment....414
VVol resource group configurations....414
Reviewing the tasks for creating block storage VVols....416
Defining a protocol endpoint for block storage VVols....417
Determining ALU and SLU volume attributes for block storage VVols....418
Creating storage containers for block storage VVols....418
Expanding the capacity of a storage container for block storage VVols....420
Changing a capability profile for block storage VVols....422
Identifying information that exists between VMs and block storage VVols....423
Deleting storage containers for block storage VVols....423

Removing a VVol environment for block storage VVols..... 424
Deleting a protocol endpoint for block storage VVols.....425
VVols and file storage..... 426
VVol configuration and operation tasks for file storage.....426
Creating storage containers for file storage VVols..... 427
Changing the VASA provider for file storage VVols..... 428
Modifying storage containers for file storage VVols..... 429
Expanding the capacity of a storage container for file storage VVols..... 429
Deleting storage containers for file storage VVols..... 430
Removing a VVol environment for file storage VVols..... 430

11 Optimizing storage performance.....433

About optimizing storage..... 434
About optimizing HBA configurations..... 435
Adding an HBA..... 436
About high temperature mode.....437
Enabling high temperature mode for VSP G1000, VSP G1500, and VSP F1500 storage systems.....438
Managing cache logical partitions..... 438
Creating a CLPR.....439
Migrating resources to and from a CLPR..... 440
Editing the settings of an existing CLPR.....441
Adjusting the cache capacity of a CLPR..... 443
Deleting a CLPR.....443
Data mobility..... 444
About data mobility.....444
Reports for logical groups displayed in the Mobility tab..... 446
Reports for DP pools displayed in the Mobility tab..... 447
Optimizing data placement in a logical group..... 449
Optimizing data placement in a DP pool.....449
Data migration..... 451
About data migration..... 451
Notes on performing data migration.....452
Conditions for data migration.....453
Migrating data for volume performance..... 459
Migrating data to a different storage system..... 460
Migrate data dialog box.....462
Prerequisites for migrating data from an SMI-S enabled storage system to another storage system..... 466
Migrating data from an SMI-S enabled storage system to another storage system 467

12 Grouping resources..... 471

Managing logical groups..... 472
About logical groups..... 472
Creating logical groups.....477
Viewing logical group reports.....478
Editing logical groups..... 479
Performing operations from a logical group.....479
Deleting logical groups.....480

Managing storage tiers..... 480
About tier-based storage management..... 481
Values to ensure acceptable tier performance..... 481
Creating tiers..... 481
Expanding a tier..... 482
Deleting tiers..... 483

13 Managing storage system information..... 485

Managing resource labels..... 486
About managing resource labels..... 486
Editing resource labels..... 487
Searching resource labels..... 487
Importing storage system resource labels..... 488
Managing WWNs by using nicknames..... 488
About managing WWNs by using nicknames..... 488
Editing a WWN nickname..... 490
Searching HCS resources..... 491
About searching HCS resources..... 491
Keyword search..... 492
Criteria search..... 493
Criteria search dialog box..... 494
Generating resource reports..... 496
About generating reports..... 497
Exporting host information to a CSV file..... 497
Exporting file server information to a CSV file..... 498
Exporting logical group information to a CSV file..... 499
Exporting search results to a CSV file..... 499
Exporting access control information for resources to a CSV file..... 500
Using reports to verify system changes..... 500
Viewing a Device Manager - Storage Navigator report..... 500
Viewing a report in the Reports window..... 501
Creating a configuration report..... 501
Deleting a configuration report..... 502

14 Managing tasks..... 503

About tasks..... 504
Customizing which tasks display in the global monitoring bar..... 507
Viewing HCS task status..... 507
Rescheduling HCS tasks waiting to be executed..... 508
Stopping running data migration or data placement profile tasks..... 508
Canceling scheduled HCS tasks..... 509
Moving HCS tasks to the HCS Task History tab..... 509
Restarting a failed or stopped HCS task..... 510
Deleting HCS tasks..... 511
Viewing system task status..... 511
Managing system tasks..... 512
Troubleshooting system tasks..... 513
Viewing data collection task status..... 514
Restarting a data collection task..... 515

15 Monitoring managed resources and resolving alerts..... 517
Monitoring storage resources..... 518
Viewing information summaries..... 519
About checking system status in the dashboard..... 519
Accessing the dashboard..... 521
Customizing the dashboard..... 522
About the Storage Systems tree..... 523
Viewing current storage system information..... 523
Viewing MP Blade information..... 524
About the Hosts tree..... 524
Viewing current host information..... 525
About the file server view..... 525
Monitoring file servers..... 525
Analyzing storage system performance..... 526
About analyzing storage system performance..... 526
Notes on health check reports..... 528
Identifying storage system performance problems..... 530
Analyzing and load balancing MP Blades or MP Units..... 532
Analyzing storage performance on a host basis..... 534
Metrics of volumes of VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, VSP F1500, Virtual Storage Platform, and HUS VM (Identify Performance Problems wizard)..... 537
Metrics of volumes of Universal Storage Platform V/VM (Identify Performance Problems wizard)..... 538
Metrics of volumes of midrange storage systems (Identify Performance Problems wizard)..... 539
Creating periodic health check reports for storage systems performance analysis 540
Editing periodic health check schedule..... 542
Exporting health check reports..... 542
Viewing and deleting health check reports..... 543
Metrics of resources of VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, VSP F1500, Virtual Storage Platform, and HUS VM (Health Check Report) 544
Metrics of resources of Universal Storage Platform V/VM (Health Check Report).. 545
Metrics of resources of midrange storage systems (Health Check Report)..... 546
Changing performance threshold settings to correctly analyze storage system performance..... 546
Configuring the Performance Reporter display..... 548
Performance Reporter and Analytics tab display differences..... 550
Analyzing Universal Replicator performance..... 552
About analyzing Universal Replicator performance..... 552
Prerequisites for analyzing Universal Replicator performance..... 553
Monitoring Universal Replicator performance..... 553
Configuring network bandwidth for analysis..... 555
Configuring metric thresholds for M-JNL analysis..... 555
Metrics for Universal Replicator performance analysis..... 556
Analyzing Universal Replicator performance in wizard mode..... 559
Analyzing Universal Replicator performance in advanced mode..... 560
Exporting Universal Replicator performance reports..... 561
About the P-VOL Write Transfer Rate metric..... 562

About refreshing UR Performance Analysis data..... 563
Managing alerts.....563
About alerts.....563
Confirming an alert..... 564

A Managing storage resources that have virtual IDs (VSP and HUS VM) 565

About managing storage resources that have virtual IDs..... 566
Displaying virtual ID information.....568
Allocating volumes to hosts with virtual IDs.....568

B Linking related products.....571

Launching other Hitachi Command Suite products..... 572
About launching other Hitachi Command Suite products..... 572
Starting related products..... 573
Starting related products from a list of hosts.....573
Notes on starting Element Manager.....573
Starting Element Manager..... 575
Starting System Management Unit (SMU)/NAS Manager..... 576
Starting File Services Manager (HFSM)..... 577
Linking with Compute Systems Manager..... 577
About linking with Compute Systems Manager..... 577
Resynchronizing hosts of Compute Systems Manager..... 578

Glossary..... 581

Index..... 595

Preface

This manual provides information for Hitachi Command Suite (HCS).

- Intended audience
- Product version
- Release notes
- Related documents
- Document conventions
- Conventions for storage capacity values
- Accessing product documentation
- Getting help
- Comments

Intended audience

This document provides software usage instructions for storage administrators.

Product version

This document revision applies to HCS v8.5.0.

Release notes

Read the release notes before installing and using this product. They may contain requirements or restrictions that are not fully described in this document or updates or corrections to this document. Release notes are available on Hitachi Data Systems Support Connect: https://knowledge.hds.com/Documents.

Related documents

Product user documentation is available on Hitachi Data Systems Support Connect: https://knowledge.hds.com/Documentation. Check this site for the most current documentation, including important updates that may have been made after the release of the product.

- *Hitachi Command Suite Installation and Configuration Guide*, MK-90HC173
- *Hitachi Command Suite Administrator Guide*, MK-90HC175
- *Hitachi Command Suite CLI Reference Guide*, MK-90HC176
- *Hitachi Command Suite Tiered Storage Manager CLI Reference Guide*, MK-90HC177
- *Hitachi Command Suite Messages*, MK-90HC178
- *Hitachi Command Suite Audit Log Reference Guide*, MK-92HC213
- *Hitachi Command Suite Mainframe Agent Installation and Configuration Guide*, MK-96HC130
- *Hitachi Command Suite System Requirements*, MK-92HC209

Document conventions

This document uses the following typographic conventions:

Convention	Description
Bold	• Indicates text in a window, including window titles, menus, menu options, buttons, fields, and labels. Example:

Convention	Description
	Click **OK**. • Indicates emphasized words in list items.
Italic	• Indicates a document title or emphasized words in text. • Indicates a variable, which is a placeholder for actual text provided by the user or for output by the system. Example: `pairdisplay -g group` (For exceptions to this convention for variables, see the entry for angle brackets.)
`Monospace`	Indicates text that is displayed on screen or entered by the user. Example: `pairdisplay -g oradb`
< > angle brackets	Indicates variables in the following scenarios: • Variables are not clearly separated from the surrounding text or from other variables. Example: `Status-<report-name><file-version>.csv` • Variables in headings.
[ ] square brackets	Indicates optional values. Example: [a | b] indicates that you can choose a, b, or nothing.
{ } braces	Indicates required or expected values. Example: { a | b } indicates that you must choose either a or b.
| vertical bar	Indicates that you have a choice between two or more options or arguments. Examples: [a | b] indicates that you can choose a, b, or nothing. { a | b } indicates that you must choose either a or b.

This document uses the following icons to draw attention to information:

Icon	Label	Description
	Note	Calls attention to important or additional information.
	Tip	Provides helpful information, guidelines, or suggestions for performing tasks more effectively.
	Caution	Warns the user of adverse conditions and/or consequences (for example, disruptive operations, data loss, or a system crash).
	WARNING	Warns the user of a hazardous situation which, if not avoided, could result in death or serious injury.

Conventions for storage capacity values

Physical storage capacity values (for example, disk drive capacity) are calculated based on the following values:

Physical capacity unit	Value
1 kilobyte (KB)	1,000 (10^3) bytes
1 megabyte (MB)	1,000 KB or $1{,}000^2$ bytes
1 gigabyte (GB)	1,000 MB or $1{,}000^3$ bytes
1 terabyte (TB)	1,000 GB or $1{,}000^4$ bytes
1 petabyte (PB)	1,000 TB or $1{,}000^5$ bytes
1 exabyte (EB)	1,000 PB or $1{,}000^6$ bytes

Logical storage capacity values (for example, logical device capacity) are calculated based on the following values:

Logical capacity unit	Value
1 block	512 bytes
1 cylinder	Mainframe: 870 KB Open-systems: • OPEN-V: 960 KB • Others: 720 KB
1 KB	1,024 (2^{10}) bytes
1 MB	1,024 KB or $1{,}024^2$ bytes
1 GB	1,024 MB or $1{,}024^3$ bytes
1 TB	1,024 GB or $1{,}024^4$ bytes
1 PB	1,024 TB or $1{,}024^5$ bytes
1 EB	1,024 PB or $1{,}024^6$ bytes

Accessing product documentation

Product user documentation is available on Hitachi Data Systems Support Connect: https://knowledge.hds.com/Documents. Check this site for the most current documentation, including important updates that may have been made after the release of the product.

Getting help

Hitachi Data Systems Support Connect is the destination for technical support of products and solutions sold by Hitachi Data Systems. To contact technical support, log on to Hitachi Data Systems Support Connect for contact information: https://support.hds.com/en_us/contact-us.html.

Hitachi Data Systems Community is a global online community for HDS customers, partners, independent software vendors, employees, and prospects. It is the destination to get answers, discover insights, and make

connections. **Join the conversation today!** Go to community.hds.com, register, and complete your profile.

Comments

Please send us your comments on this document to doc.comments@hds.com. Include the document title and number, including the revision level (for example, -07), and refer to specific sections and paragraphs whenever possible. All comments become the property of Hitachi Data Systems Corporation.

Thank you!

1

Overview of Hitachi Command Suite

Hitachi Command Suite (HCS) is a comprehensive software suite providing management services for storage systems and hosts. Storage configuration, virtualization, reporting, and monitoring tools are fully supported.

- ☐ About Hitachi Command Suite
- ☐ Features
- ☐ What's new
- ☐ System configuration
- ☐ Process flow
- ☐ Navigating the interface
- ☐ Navigating help

About Hitachi Command Suite

HCS consists of a number of storage management software products used for managing storage resources in large-scale, complex SAN environments.

HCS includes:

- **Hitachi Device Manager (HDvM):** Supports registration of physical resources to be managed, including storage systems, file-servers, hosts, and related tasks such as volume allocation and grouping of resources for easier management and access.
- **Hitachi Tiered Storage Manager (HTSM):** Supports storage tiers of differing performance characteristics so that volume data storage costs and performance can be optimized.
- **Hitachi Replication Manager (HRpM):** Supports volume data replication for backup and disaster recovery.
- **Hitachi Tuning Manager (HTnM):** Supports optimizing the performance of storage resources.
- **Hitachi Compute Systems Manager (HCSM):** Supports centralized monitoring and management of hosts, including rebooting and power management.
- **Hitachi Command Director (HCmD):** Supports sophisticated business centric views of storage environments, such as adherence to service level agreements and data usage trends to forecast future storage requirements.
- **Hitachi Automation Director (HAD):** Supports automation of storage management and provisioning tasks through the use of pre-configured service templates.
- **Hitachi Dynamic Link Manager (HDLM):** Supports the use of multiple paths between resources such as hosts and storage for path fail-over and load balancing.
- **Hitachi Global Link Manager (HGLM):** Supports management of multi-path management software between resources, such as hosts and storage.

Each product must be licensed for use in HCS. At a minimum, you must license Device Manager. Additional licensing can be added as needed for other storage management products. Related functionality becomes available in the HCS user interface in the form of activated menu choices, and new or updated tabs and related screens and buttons.

For Hitachi Virtual Storage Platform F400, F600, F800 (VSP Fx00 models), Hitachi Virtual Storage Platform G200, G400, G600, G800 (VSP Gx00 models), Hitachi Virtual Storage Platform G1000 (VSP G1000), Hitachi Virtual Storage Platform G1500 (VSP G1500), and Hitachi Virtual Storage Platform F1500 (VSP F1500), additional functions for configuring the storage system are available from launch points in HCS to Hitachi Device Manager - Storage Navigator (HDvM - SN).

Instances of the management software products above can be installed on one or more HCS servers depending on the scale of resources under management and geographic location. For information about supported HCS server operating systems and other system requirements, see *Hitachi Command Suite System Requirements*. For information about product installation options, see the *Hitachi Command Suite Installation and Configuration Guide.*

Related concepts

- Features on page 23
- What's new on page 24
- System configuration on page 25
- Navigating the interface on page 29
- Navigating help on page 31

Related tasks

- Logging in to Hitachi Command Suite on page 41

Features

Hitachi Command Suite provides a wide variety of security, scalability, data migration, replication, performance, and administrative features for managing your storage system needs in this release.

Hitachi Command Suite includes the following functionality:

- Server architecture supports resource scalability (millions of storage objects) and geographic scalability.
- Provides secure remote management over Internet and wide area networks (WANs) using a sophisticated web client.
- Provides multiple levels of security and access for storage administrators, integration with external authentication servers, and use of resource groups to control access to specific resources.
- Security tools to prevent unauthorized access, such as account locking, and restricted login retry attempts.
- Provides for automation scripts using component CLIs.
- Supports agentless discovery and mapping of servers (hosts) and storage.
- Supports hosts such as Microsoft Windows, Sun Solaris, HP-UX, IBM AIX, and Linux.
- Supports virtualization servers such as Microsoft Hyper-V and VMware ESXi.
- Supports FC, iSCSI, and FCoE connected hosts.
- Supports a wide variety of block, file, unified (hybrid) and SMI-S compliant storage systems.
- Supports virtualization of volumes from external storage.
- Supports volume migration between supported source and target storage systems.

- Provides simplified volume provisioning in complex environments.
- Provides logical, physical, and host view storage management.
- Provides efficient, cost effective use of storage capacity as needed using DP pools.
- Supports optimized application performance with an installed Tiered Storage Manager license when DP pools are configured to support performance tiers. Using the Mobility tab, you can optimize data placement tasks when migrating online data between tiers.
- Supports analyzing storage system performance with an installed Tuning Manager license, by using the Analytics tab.
- Integrates with other products and software tools that provide extended capabilities. For example, integration with Hitachi Command Director supports SLA/SLO compliance monitoring and reporting and links with native storage management tools where necessary.
- Provides volume replication services for data protection and supports analyzing Universal Replicator C/T delta performance on the Replication tab (with an installed Replication Manager and Tuning Manager license).
- Provides keyword and criteria based searching of managed resources, and managed resource reporting with data export to CSV files.

Related concepts

- About Hitachi Command Suite on page 22

What's new

Hitachi Command Suite v8.5.0 includes new and enhanced functions.

The following list describes a few of the key updates in this version:

- The Virtual Storage Platform family has been expanded to include the VSP G1500 and VSP F1500 models.
- HCS for VVols (block storage environment) has been improved so that the number of available LDEV IDs can be identified in the GUI, and container IDs are also displayed to help your site identify the storage containers that share the same IDs between the vCenter client and HCS.
- GAD Quorum for active/active has been updated so that your site can set a time-out value for a health check with the DKC.
- The procedure workflow (such as wizard messages, alerts and labels within the GUI) have been improved within the Manually Change Task Status dialog and Change Host I/O paths window.
- The Replication tab can be used to expand or shrink the copy topology by changing the type.
- As part of volume creation, you can allocate the deduplication system data volume (DSD volume) to a DP pool or create a DP volume for which the capacity saving function (dedupe and compression) is enabled.
- The capacity saving function (dedupe and compression) is supported for flash-based media types: FMD, FMD DC2, SSD and FMD HD.

This list highlights new software (GUI) features. It is not an exhaustive list of enhancements.

For complete information about new features and enhancements, see *Hitachi Device Manager Release Notes* and *Hitachi Tiered Storage Manager Release Notes*.

For a complete list of system requirements for the management server, Host Data Collector, Device Manager agent, CLI, and storage systems, see *Hitachi Command Suite System Requirements*.

Related concepts

- About changing the copy topology type on page 307
- Working with VMware VVols on page 406

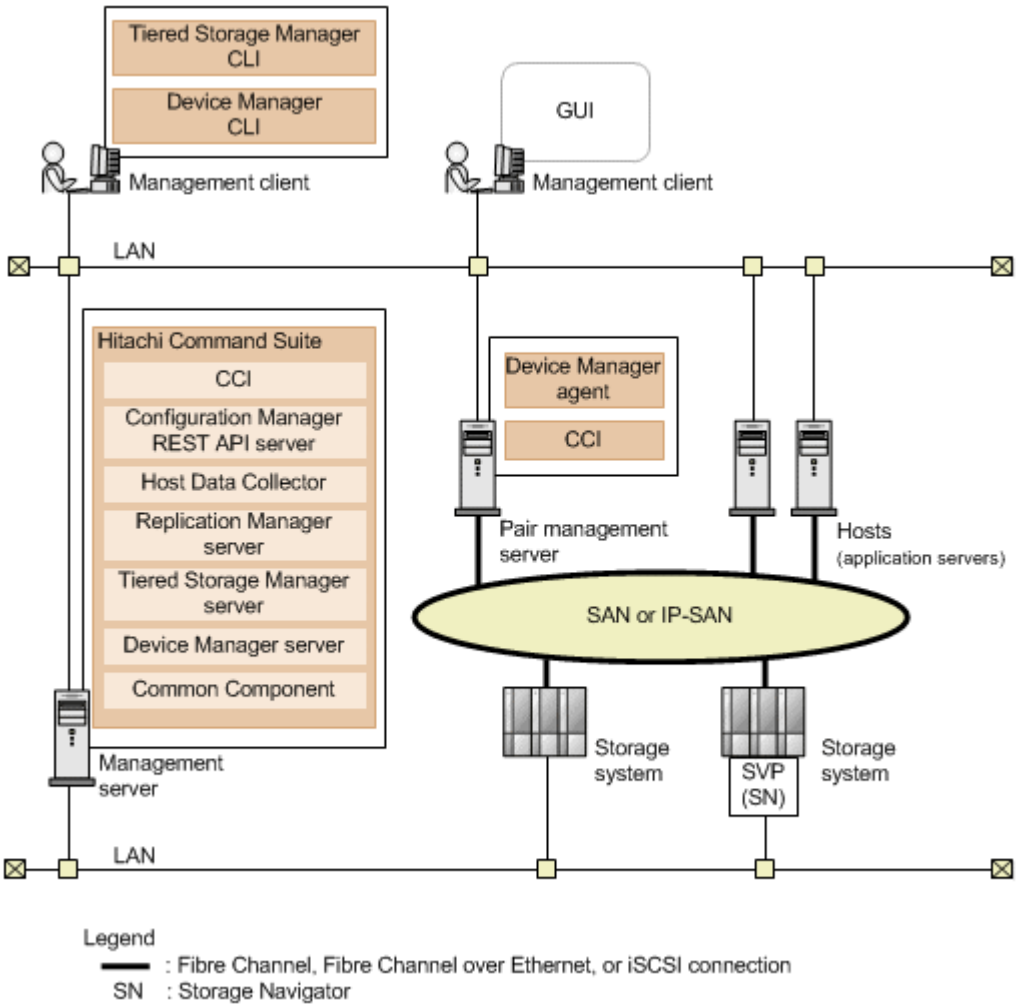
Related tasks

- Configuring the quorum disk on page 345
- Identifying information that exists between VMs and block storage VVols on page 423
- Determining ALU and SLU volume attributes for block storage VVols on page 418

System configuration

As part of a basic storage network configuration, Hitachi Command Suite is comprised of HCS server-based components, agents, and clients that enable you to manage storage resources.

The following figure illustrates Hitachi Command Suite management server components, and the basic configuration of the storage network.



Management servers

- The management server is where HCS is installed.
- The management server communicates with management clients, storage systems, pair management servers, file-servers, and hosts over LAN connections.
- Additional software can be installed on the management server to provide extended management capabilities.
- The management server can be configured in an active-standby cluster configuration consisting of two physical servers.

Management server components

The HCS base installation consists of the following components, which are always installed or removed together on the management server:

- Hitachi Command Suite Common Component
  Provides user account management, security monitoring, and other functions common to all HCS products.
- Device Manager server

HCS uses this component to manage storage system volumes.

- Tiered Storage Manager server
  Tiered Storage Manager uses this component to manage storage system volume migration.
- Replication Manager server
  Replication Manager uses this component to manage storage system volume replication.
- Host Data Collector
  HCS uses this component to collect information about the volumes used by the hosts.



Note: The Host Data Collector component can be installed on other servers and accessed remotely by HCS.

Management clients

- Manage storage resources from the management server by using the HCS web client (GUI), or by using the CLI client software to issue commands.
- CLI client components (Device Manager and Tiered Storage Manager) require a separate installation from the web client.
- From the web client, with a Device Manager license, launch Replication Manager to use a subset of Replication Manager functionality.

Hosts

- File servers and hosts (application servers) access volumes in storage systems that are managed by HCS.
- Hosts access storage over SAN (Fibre Channel) or LAN (iSCSI and FCoE) connections.
- Hosts can be virtualization servers (VMware ESXi) and their virtual machines, and mainframe hosts.
- Hosts can be file servers (NAS Platform, Hitachi NAS Platform F (HNAS F), and Hitachi Data Ingestor (HDI)).

File servers and hosts (application servers) access volumes in storage systems that are managed by HCS over a Storage Area Network (SAN) or an IP-SAN.

Pair management servers

- Collects management information about copy pair configurations and related status information, and provides for copy pair operations.
- Command Control Interface and Device Manager agent are installed for copy pair monitoring and management.

For information about performing the base installation, see the *Hitachi Command Suite Installation and Configuration Guide*.

For information about customizing and extending the base installation, see the *Hitachi Command Suite Administrator Guide*.

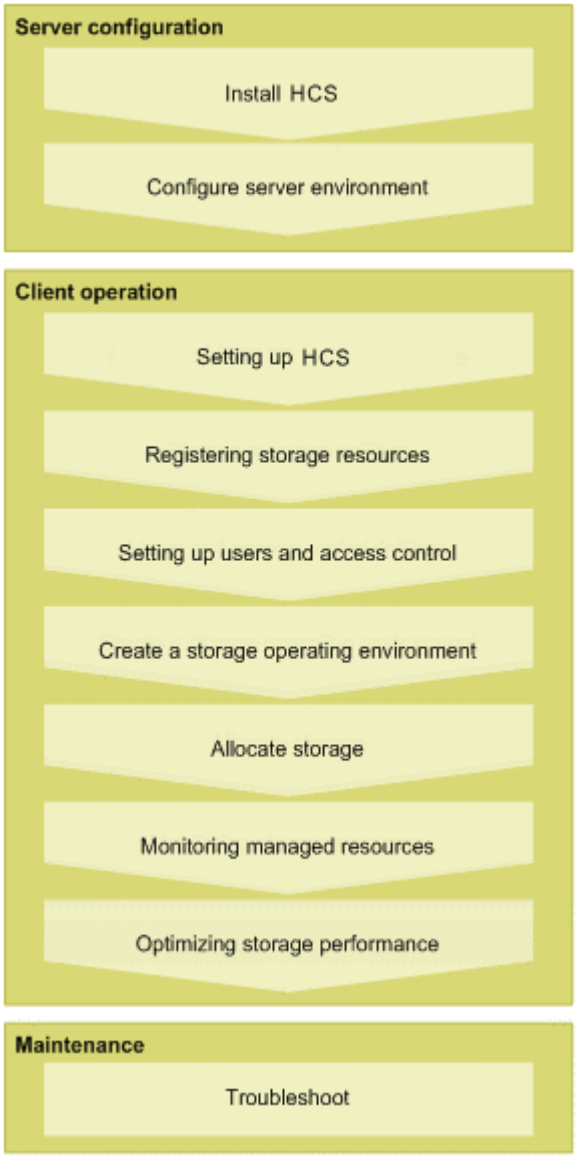
Related concepts

- About Hitachi Command Suite on page 22

Process flow

The following graphic illustrates the flow of system operations when using Hitachi Command Suite and Hitachi Tiered Storage Manager.

- Server configuration steps are related to installing and configuring the server itself.
  For more information on server installation and configuration, see:
  - *Hitachi Command Suite Installation and Configuration Guide*
  - *Hitachi Command Suite Administrator Guide*
- Client operation steps are illustrated and explained further with workflow graphics and comments in this user guide.
- Maintenance refers to troubleshooting the server if issues arise when running Hitachi Command Suite.



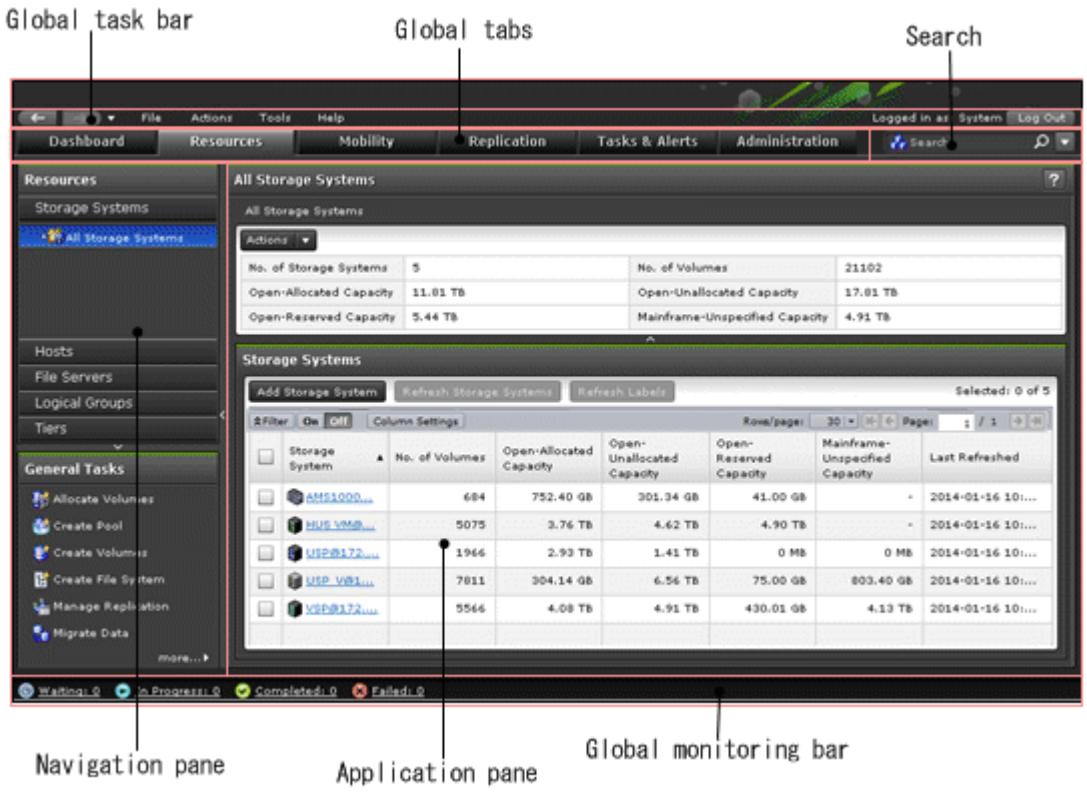
Navigating the interface

Hitachi Command Suite provides a sophisticated interface with elements such as menus, tabs, navigation options, application details, status information and search options supporting ease of use and flexibility for storage and server resources.

Interface elements

- Global task bar - always visible, indicates the date and time of the last log on event, forward/back navigation buttons, clickable navigation history (down arrow), menu access to licensing and help. (Help), launching licenced options (Tools), and Log Out for exiting Hitachi Command Suite.
- Global tabs - always visible, provides access to applications.

- Search - always visible, provides keyword and criteria-based search.
- Navigation pane - differs by tab, provides access to resources and commonly-used tasks.
- Application pane - differs by tab, provides resource summary information, resource object list, and related drill down details.
- Global monitoring bar - always visible, provides links for submitted tasks by status.



Navigation pane

- Accordion menus provide easy access to resource trees.
- Resize panes by dragging divider bars. Collapse panes by clicking divider bar arrows.
- Access frequently needed tasks under General Tasks. Click More... to see hidden general tasks.

Application pane

- Minimize and expand panes by using the double arrow symbol in the title.
- Click Column Settings to display or hide columns, change column order, or view column descriptions.

- Right-click a table heading and select menu options, such as Hide Column or Show all Columns.
- Arrange columns by using drag-and-drop.
- Sort lists by clicking the column title.
- Navigate large lists by using Page controls.
- Click Filter to reduce large lists, or to find specific items by defining specific search conditions. Filter allows multiple conditions to be defined.
- In a list, click a link to display more detail about the item. As you drill-down, the breadcrumb list (*item* > *detail*) above the summary pane is updated and serves as a useful navigation tool.
- In a list, rows are highlighted as you roll your mouse over them, indicating your row position. To select a specific item, select the desired check box or click the row and the darker highlight indicates the row is selected.
- To select multiple rows, select the desired check boxes or rows. You can also use Shift+click to select a range of items. To select all rows, select the check box for the title row, or clear the check box to de-select all rows.
- Selecting rows implies you intend to perform an action on the selected item or items. Actions are initiated with buttons or from the Actions menu.
- To copy cell or row data, select one or more rows with data, right-click and select Copy This Cell or Copy Selected Rows. This is useful for emailing small amounts of data about a storage resource. If you select empty rows, the copy options do not appear when you right-click. For reporting on large numbers of objects and for more complete data, use CSV export.

Related concepts

- About Hitachi Command Suite on page 22
- Navigating help on page 31

Navigating help

The Help system provides brief explanations of the features of this product and helps you understand its capabilities. Navigating is the means by which you access the information in the Help system.

When you access Help, the navigation pane displays.

Navigating

- To navigate among topics, use the navigation pane, or click Previous topic or Next topic in the button bar.

- Use the breadcrumb trail at the top of each topic to see your location, or to return to a higher level topic.
- To find information for a specific topic, click the Related topics links.

Using the navigation pane

- Contents
  Displays topics and sub-topics. As you move through Help, the current topic is highlighted.
- Search
  Displays search results. Click the entries to view contents.

Printing topics

- To print topics, right-click the topic and select Print or click the printer icon on the button bar.

Related concepts

- About Hitachi Command Suite on page 22
- Navigating the interface on page 29

2

Setting up Hitachi Command Suite

This module describes how to configure basic Hitachi Command Suite settings.

- ☐ Configuring Hitachi Command Suite
- ☐ Configuring your browser and Java for Hitachi Command Suite
- ☐ Logging in to Hitachi Command Suite
- ☐ Logging in when HCS is not available
- ☐ Setting up security
- ☐ Downloading components
- ☐ Managing HCS licenses
- ☐ Managing licenses for Virtual Storage Platform storage systems

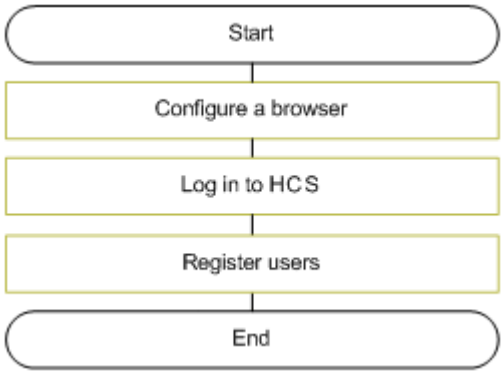
Configuring Hitachi Command Suite

Initial configuration of HCS, including configuring access to Hitachi Device Manager - Storage Navigator requires configuring a browser, licensing, logging in, and registering users.

The following graphic illustrates the required tasks:

- Configure your browser to correctly display and run Hitachi Command Suite. If you have issues, verify your browser setup.
- Log in to HCS. Note that you can register licenses from the login screen.
- Register your users (create user accounts).

Note: Typically, user accounts are created so that storage administration can be shared or delegated, or to control access to specific storage resources.



Related tasks

- Creating a user account on page 91

Configuring your browser and Java for Hitachi Command Suite

To communicate with the Hitachi Command Suite management server and Hitachi Device Manager - Storage Navigator or the maintenance utility, you must configure your browser and Java settings.

About configuring browser settings

Before using a browser to access Hitachi Command Suite or Device Manager - Storage Navigator, configure security and other settings.

Perform the following configuration tasks:

Note: For VSP G1000, VSP G1500, and VSP F1500, configure the IP address or host name of the service processor (SVP). For VSP Fx00 models and VSP

Gx00 models, configure the IP addresses of the SVP and the maintenance utility.

- Disable pop-up blocking
- Disable plug-ins
- Set security options
- Configure proxy settings
- Configure log output settings
- Configure Java Web Start settings
- Clear your browser's cache when upgrading
- Browser must allow first-party, third-party, and session cookies

Note: For specific instructions on configuring your browser, refer to the browser product documentation.

After you configure your browser settings, verify the following:

- The Java™ software environment is installed and the Java software is configured
- Communications with HCS are secure
- Check the management server name resolution

Note: If you have issues with general access, or with using a component of Hitachi Command Suite, first verify that the problem is not related to the management server, or a network connectivity issue.

For troubleshooting problems related to Device Manager servers, see the *Hitachi Command Suite Administrator Guide*.

Related tasks

- Checking management server name resolution on page 35
- Setting options for Internet Explorer on page 37
- Setting options for Firefox on page 38
- Setting the Java™ Web Start (versions 6.0 and 7.0) proxy to link with other products on page 39
- Configuring JRE versions from JWS version 6.0 on page 40
- Clearing the cache when upgrading Hitachi Command Suite on page 41

Checking management server name resolution

HCS management servers host names must resolve to an IP address.

After logging in to Device Manager, before you can start another HCS product, such as Replication Manager or Tuning Manager, the host name of the management server must be specified so that it can be resolved to the IP address. If the name cannot be resolved, startup of the HCS product might fail.

Note: If you are having issues with general access, or with using a component of Hitachi Command Suite, first make sure the problem is not related to the management server, or related to networking.

For troubleshooting problems related to Device Manager servers, see the *Hitachi Command Suite Administrator Guide*.

Procedure

1. Launch your browser.
2. Verify whether you can access the management server on which the products such as the Device Manager server, Tuning Manager server, or Replication Manager server are installed: Use the 'ping' command to specify the host name of the target server.

Related concepts

- About configuring browser settings on page 34

Disabling pop-up blocking

You can disable pop-up blocking, or configure URLs, so that HCS can display all windows.

If you are using HCS on a browser for which pop-up blocking is enabled, you must disable pop-up blocking.

Alternatively, you can register the URLs of HCS products and the IP address or host name of SVP as allowed sites. For VSP Fx00 models and VSP Gx00 models, to start the maintenance utility, register the IP addresses of both controller 1 and controller 2 of the storage system.

If SSL or TLS is being used for communication between the management server and the management client, also register the URLs for SSL communication in the pop-up blocker settings.

Before you begin

Obtain the IP address or host name of the Device Manager server.

Procedure

1. Launch the browser.
2. To permit the display of pop-ups, add the following URLs to either allowed websites or to the list of sites where popup controls are excepted, depending on your browser's settings:

   For SSL: `https://`*`IP-address-or-host-name-of-the-Device-Manager-server`*

   For non-SSL: `http://`*`IP-address-or-host-name-of-the-Device-Manager-server`*

Related concepts

- About configuring browser settings on page 34

Setting options for Internet Explorer

To communicate with the HCS management server and for correctly displaying windows when using Internet Explorer, configure Internet Explorer options.

Procedure

1. Start Internet Explorer.
2. From the **Tools** menu, select **Internet Options**.
3. In the **Internet Options** dialog box, click the **Security** tab, and then click **Trusted sites**.
4. Click **Sites** and add the IP address or host name (alias name) for Device Manager and the related management servers that contain the software to be started from Device Manager. To manage VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, or VSP F1500, you must also add the URL for the SVP. For VSP Gx00 models, to start the maintenance utility, register the IP addresses of both controller 1 and controller 2 of the storage system.
5. Click **Custom Level**.
6. In the **Security Settings** dialog box, verify that the items in the dialog box are configured as follows:
   - Run ActiveX controls and plug-ins is set to **Enable**.
   - Script ActiveX controls marked safe for scripting is set to **Enable**.
   - Launching programs and files in an IFRAME is set to **Prompt** or **Enable**.
   - Submit non-encrypted form data is set to **Prompt** or **Enable**.
   - Active scripting is set to **Enable**.
   - File download is set to **Enable**.
   - Allow Web sites to prompt for information using scripted windows is set to **Enable**.
7. In the **Internet Options** dialog box, select the **Privacy** tab, and then enable cookies.
8. In the **Internet Options** dialog box, select the **Advanced** tab, then select **Warn about certificate address mismatch**, and **Show pictures**. For VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, or VSP F1500, clear the **do not save encrypted pages to disk check-box**.
9. If the logged-in user does not have administrator permissions and uses Internet Explorer 10 or 11, disable the protected mode.
10. From the **Tools** menu, select **Manage Add-ons**, and enable **Shockwave Flash Object**.

11. If Internet Explorer 10 or 11 is used in Windows Server 2012, start the server manager. Select in the sequence of **Dashboard**, **Add roles and features** wizard, and **Features**, and then install **Desktop Experience**.
12. When using Device Manager for the first time with Windows Server 2008, in the **Internet Options** dialog box, select the **Advanced** tab, and then clear the **Do not save encrypted pages to disk** check box.

Result

Settings are applied.

Note that when using Internet Explorer in Windows Server 2008 or Windows Server 2012, if the loading animation does not start or if files cannot be downloaded from a server for which HTTPS is enabled, disable Internet Explorer Enhanced Security Configuration.

If you use Internet Explorer 10 or 11, specify the same protected mode for the Internet domain for accessing Device Manager and for the Internet domains of other programs that cooperative with Device Manager. If different modes are specified, the other programs cannot start. For example, in the case that Device Manager and Tuning Manager server are connected remotely.

The following problems might occur:

- Online Help might not display properly.
- An error message displays that indicates that Adobe Flash Player is not installed.

If you encounter these problems, disable ActiveX filtering, and from Internet Explorer, in Compatibility View settings, register the IP address or the host name of the Device Manager server. If the online Help does not display properly even after these settings are applied, press F5 to refresh the browser window.

Tip: Set the text size to Medium or Larger. If you set text size to Largest, text characters might overlap.

Related concepts

- About configuring browser settings on page 34

Related tasks

- Disabling pop-up blocking on page 36

Setting options for Firefox

To communicate with the HCS management server, and for correctly displaying windows when using Firefox, configure Firefox options.

Procedure

1. Start Firefox.
2. In the environment settings window, set the items as follows:
   - Enable first-party, third-party, and session cookies.
   - Enable JavaScript.
   - Use the default font.

Related concepts

- About configuring browser settings on page 34

Setting options for Google Chrome

To communicate with the HCS management server, and for correctly displaying windows when using Google Chrome, configure Google Chrome options.

Procedure

1. Start Google Chrome.
2. From the menu, select **Settings** then click **Show advanced settings...**. In the **Privacy** section, click **Content settings...**
3. Set the following items:
   a. To enable cookies, in the **Cookies** section, select **Allow local data to be set (recommended)**.
   b. To enable JavaScript, in the **JavaScript** section, select **Allow all sites to run JavaScript (recommended)**.

Specifying settings to link with Device Manager - Storage Navigator

To manage VSP G1000, VSP G1500, VSP F1500, VSP Gx00 models, or VSP Fx00 models by using Hitachi Command Suite (HCS), enable TLS settings to link with Device Manager - Storage Navigator.

Procedure

Enable TLS settings for the browser.

Setting the Java™ Web Start (versions 6.0 and 7.0) proxy to link with other products

You must set up Java Web Start so that it uses the correct version of JRE to execute the JAVA GUI to run with other products. This procedure is for Java Web Start versions 6.0 and 7.0.

Note: To connect to Device Manager - Storage Navigator behind restrictive firewalls, you must set up a proxy to route communications between HCS and Device Manager - Storage Navigator, and then configure Java Web Start to use the correct proxy settings to enable communication.

Procedure

1. Start the Java Control Panel.
2. On the **General** tab, click **Network Settings**.
3. In the **Network Settings** window, select the method to be used to set up the proxy.
   - If the proxy setting is enabled in the web browser, select **Use browser settings**, and then proceed to step 7.
   - To manually set up a proxy, select **Use proxy server**, enter the following values, and proceed to step 4.
     - **Address**: IP address or name of the proxy server
     - **Port**: Port number of the proxy server
   - If not setting a proxy, select **Direct connection** and proceed to step 7.
4. Click **Advanced**.
   The **Advanced Network Settings** window appears.
5. In the exceptions field, enter the IP address or the name of the Device Manager server.
6. Click **OK**.
7. In the **Network Settings** window, click **OK**.
8. In the **Java Control Panel**, click **OK** and exit the control panel.

Related concepts

- About configuring browser settings on page 34

Configuring JRE versions from JWS version 6.0

When there are multiple versions of JRE and some of these versions are the same as or later than the version required for a Java application, Java™ Web Start determines which version of JRE to use.

You must set up Java Web Start version 6.0 so that it uses the correct version of JRE to start the GUI.

Note: For GUI requirements, see the *Hitachi Command Suite System Requirements*.

Procedure

Set up the JRE version required to use Java control panel. Use the setup method appropriate to the JWS version.

- For JWS version 6.0, use the Java control panel.
- For GUI requirements, see the *Hitachi Command Suite System Requirements*.

Related concepts

- About configuring browser settings on page 34

Clearing the cache when upgrading Hitachi Command Suite

After upgrading HCS, you must clear the cache for Java™ Web Start and for your browser.

Procedure

1. For Java™ Web Start version 6.0 use the Java Control Panel to delete temporary internet files.
2. For the browser you are using, delete temporary internet files.

Related concepts

- About configuring browser settings on page 34

Logging in to Hitachi Command Suite

To log in to Hitachi Command Suite use the correct URL in your browser and register licenses if necessary.

Before you begin

Before you can log in to Hitachi Command Suite, you must first register valid licenses using the Licenses button.

Note: To log in with a user account other than a built-in user account, you must first set up new user accounts, or user groups and roles. This requires the User Management permission, and the Admin role for All Resources.

Procedure

1. Start a web browser and enter the URL of the HCS server:
   `http://`*`server-IP-or-name`*`:`*`port-number`*`/DeviceManager/`
   - *server-IP-or-name*: IP address or host name of the Device Manager server.
   - *port-number*: Port number of the HBase Storage Mgmt Web Service.

   URL examples:
   - `http://localhost:22015/DeviceManager/`
   - `https://localhost:22016/DeviceManager/`

   Note: Using the localhost name entry implies you are logged in at that server. To access the interface from another system, specify an IP address or host name in the URL. For a secure connection, use the https URL on port 22016.

2. Enter values in the **User ID** and **Password** fields, and click **Log In**.

Tip: To authenticate with an external authentication server, use its authentication password. When an external authorization server is linked, use the user ID and password that are registered in the external authorization server.

Related concepts

- About HCS license management on page 46
- About user accounts and controlling access to resources on page 88
- About access control on page 99
- About launching other Hitachi Command Suite products on page 572
- About Hitachi Command Suite on page 22
- About configuring browser settings on page 34

Related tasks

- Registering an HCS license on page 47

Related references

- User ID and password policies on page 92

Logging in when HCS is not available

For the VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, and VSP F1500 storage systems, if HCS is not available and you have an administrator login account with the required permissions, you can log in directly to Device Manager - Storage Navigator.

Before you begin

- You are managing VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, or VSP F1500 storage systems.
- You must have an administrator login account with the required permissions. For information about creating user accounts in Hitachi Device Manager - Storage Navigator, see the *System Administrator Guide* for your storage system.

Procedure

1. Start a web browser.
2. Enter the URL:
   - For VSP G1000, VSP G1500, or VSP F1500 storage systems, enter: `https://IP-address-or-host-name-of-the-SVP/sanproject/emergency.do`

- For VSP G200 storage systems, enter: `https://IP-address-or-host-name-of-the-SVP/dev/storage/8320004XXXXX/emergency.do` (where the model number is '8320004' and '4XXXXX' indicates the system serial number)
- For VSP G400, G600 storage systems and VSP F400, F600 storage systems, enter: `https://IP-address-or-host-name-of-the-SVP/dev/storage/8340004XXXXX/emergency.do` (where the model number is '8340004' and '4XXXXX' indicates the system serial number)
- For VSP G800 and VSP F800 storage systems, enter: `https://IP-address-or-host-name-of-the-SVP/dev/storage/8360004XXXXX/emergency.do` (where the model number is '8360004' and '4XXXXX' indicates the system serial number)

3. The following actions might be required to open the login dialog box, depending on your environment:
   - If a message indicates that the enhanced security configuration is enabled on the computer, select **In the future, do not show this message** and click **OK**.
   - If the SVP is set to support SSL-encrypted communication and security messages appear, make sure the certificate is correct and follow the instructions in the dialog box.
   - If a messages indicates that certain web sites are blocked, make sure you have added the SVP to the trusted sites zone.
4. Enter the user ID and password for the account.
5. Click **Log In**.
6. If the Security Information dialog box appears, click **Yes**.
7. If an Adobe Flash Player local storage area pop-up dialog box appears, click **Allow** to open the Device Manager - Storage Navigator main window. The cache function of Flash Player optimizes the process of Device Manager - Storage Navigator. Denying the request might reduce processing speed.

Result

You are successfully logged in to Device Manager - Storage Navigator.

Note: If the login process fails three times by using the same user ID, Device Manager - Storage Navigator will stop responding for one minute. This is for security purposes and is not a system failure. Wait, and then try again.

Setting up security

This module describes how to configure security settings and policies.

About configuring security options

For tighter login security, you can specify security options from HCS, such as:

- Password policy conditions to prevent users from specifying easy-to-guess passwords.
- Automatic locking of user accounts when successive login attempts have failed.
- Display of a specific message (a warning banner) during user login.

You can also set security options from the management server. For details, see the *Hitachi Command Suite Administrator Guide*.

Related tasks

- Setting a password policy on page 44
- Setting automatic account locking on page 45
- Setting a warning banner message on page 45

Setting a password policy

Create a password policy to enforce the use of strong passwords.

Procedure

1. On the **Administration** tab, select **Security** to launch the security window.
2. On the **Security** window, select **Password** to display the current password policy.
3. Click **Edit Settings** and configure a password policy. For example, configure a minimum password length, and a minimum number of required uppercase, lowercase, numeric, and symbolic characters.
4. Click **OK** to save the new password policy.
5. Select **Password** again to confirm the updated password policy.

Related concepts

- About configuring security options on page 44

Related tasks

- Setting automatic account locking on page 45

Related references

- User ID and password policies on page 92

Setting automatic account locking

A user account can be automatically locked after a specified number of failed login attempts. You can specify the number of failed attempts before a user account is locked by configuring an account lock policy.

Procedure

1. On the **Administration** tab, select **Security**.
2. From the tree, select **Account Lock**.
3. Click **Edit Settings** and specify a number.
4. Confirm the settings were changed by selecting **Account Lock** in the **Security** tree and viewing the current settings.

Related concepts

- About configuring security options on page 44

Related tasks

- Setting a password policy on page 44
- Setting a warning banner message on page 45

Setting a warning banner message

You can set a warning banner message that will appear in the Login window when a user logs on.

Procedure

1. On the **Administration** tab, select **Security**.
2. From the tree, select **Warning Banner**.
3. Click **Edit Message** and enter the warning message text in the **Message** box.

   You can preview the message by clicking **Preview** and viewing the message in the **Preview** box.
4. Click **OK** to save the message.
5. Confirm the warning banner displays in the Login window.

Related concepts

- About configuring security options on page 44

Related tasks

- Setting automatic account locking on page 45
- Setting a password policy on page 44

Downloading components

This module describes how to download components.

About downloading components

A download menu allows you to download Device Manager agent, the CLI application, or Host Data Collector installation files.

Related tasks

- Downloading agents, CLI, and Host Data Collector files on page 46

Downloading agents, CLI, and Host Data Collector files

Use the download feature to download agents, the CLI application, or Host Data Collector installation files:

Procedure

1. From the **Tools** menu, select **Download**.
   The download dialog box opens.
2. Choose from the links in the dialog box to download the desired installation files.
3. Click the **ReadMe** links for installation instructions.

Related concepts

- About downloading components on page 46

Managing HCS licenses

This module describes how to manage HCS licenses.

About HCS license management

Before you can log on to Hitachi Command Suite, you must register valid HCS licenses.

Each product managed from HCS also requires a registered license.

Licenses that are registered while you are logged in are not enabled until you log out and log in again.

The following license types are available:

- Permanent
- Term
- Subscription
- Emergency

- Temporary

Note: Before you log in for the first time, a License Alert appears on the login window.

When a problem occurs, such as when a license expires or capacity limits are exceeded, an alert is displayed in the main window.

Related tasks

- Registering an HCS license on page 47
- Checking HCS license information on page 48

Registering an HCS license

Before you can log on to Hitachi Command Suite, you must register valid HCS licenses.

Each product managed from HCS also requires a registered license.

Before you begin

If you are using Firefox, before you register the license key file, disable SSL communication if SSL communication is being used between the Device Manager server and the management client. For information about changing SSL communication settings, see the *Hitachi Command Suite Administrator Guide*.

To use the functionality to manage VVols for VMware vSphere 6.0 or later, a VVol management server is required. (This is in addition to the existing management server. The VVol management server is the server on which an instance of HCS is installed, and which is specifically designated to manage VVols.) Register a license for Device Manager in HCS that is installed on the VVol management server. Do not register other licenses for HCS.

For information about system installation and configuration, see the *Hitachi Command Suite Administrator Guide.*

Procedure

1. On the login window, click **Licenses**.
   If you are already logged in, from the **Help** menu, select **About**.
2. See **License Type** and **License Messages** to review the current status of licenses.
3. Register one or more licenses using one of these methods:
   - Enter the license key manually
   - Specify the license key file (recommended)
4. Click **Save**.
   If you registered the license after you logged in, you must log out and then log in again for the license to be enabled.

5. Confirm that the license is successfully registered by viewing the displayed message.

Related concepts

- About HCS license management on page 46

Related tasks

- Checking HCS license information on page 48

Checking HCS license information

You can check license information and the license expiration date. All users can perform this task.

Procedure

1. On the login window click **Licenses**. If you are already logged in, from the **Help** menu, select **About**.
2. Check the license status displayed in **License Type** and **License Messages**.
3. Click the license name link for each product to check the storage capacity limit and expiration date for each license.

Related concepts

- About HCS license management on page 46

Related tasks

- Registering an HCS license on page 47

Managing licenses for Virtual Storage Platform storage systems

By launching the License Keys menu from the Resources tab, your site can manage Device Manager - Storage Navigator licenses for VSP G1000, VSP G1500, and VSP F1500. You can also manage licenses for VSP Fx00 models and VSP Gx00 models.

License key types

To use software, you must install the license key provided when you purchase that software.

You can use software with licensed capacity for a term key by installing a term key and overwriting a permanent key as long as the term key is valid. If the term key expires when the system is being used, and the capacity needed for the operation is insufficient, operations that you can perform are limited. In this case, a SIM that indicates the term key expiration (reference code 7ff7*xx*) is output on the Alerts tab in the Storage Systems window.

The following table describes the four types of license keys.

Type	Description	Effective term[1]	Estimating licensed capacity
Permanent	For purchase	No limit	Required
Term	For purchase	365 days	Required
Temporary	For trial use before purchase (try and buy)	120 days	Not required
Emergency	For emergency use	30 days	Not required

Notes:

1. When you log in to Device Manager - Storage Navigator, a warning message appears if 45 days or less remain before the expiration.

Related tasks

- Enabling a license on page 50
- Examples of license information on page 52
- Installing the software on page 49

Installing the software

Follow the appropriate steps for your storage system.

Procedure

1. On the **Resources** tab, click **Storage Systems**, and then expand **All Storage Systems** and the target storage system.
2. Click **License Keys**.
3. In the **License Keys** window, click **Install Licenses**.
4. Select whether to enter a key code or specify a license key file.
   - **License Key Code** or **Key Code**: Enter a key code to install the software. In **License Key Code** or **Key Code**, enter the license key code for the software.
   - **License Key File** or **File**: Specify a license key file to install the software. Click **Browse** and specify the license key file. You can use a file name of up to 200 alphanumeric characters, excluding several symbols (" \ ; : * ? < > | / ,). The file extension is "plk".
5. Click **Add** or **Apply**.
6. Click **OK** in the confirm window.

   If a software installation fails, the **Error Message** window opens. To display the cause of error, from the **Error Message** window, select the software and click **Detail**.
7. In the **Selected License Keys** table, set the status of license keys for each software application.
   - **Enable Licenses**: Installs license keys in enabled status. You can select more than one software application to install licenses for.

- **Disable Licenses**: Installs license keys in disabled status. You can select more than one software application to install licenses for.
- **Clear All**: Delete all license keys from the Selected License Keys table.

8. Click **Finish**. The **Confirm** window opens.
9. In the **Confirm** window, check the settings. In the **Task Name** field, enter a task name.
10. Click **Apply**. The task is registered. If the **Go to tasks window for status** check box is checked, the **Task** window opens.

    If a software installation fails, the **Error Message** window opens. To display the cause of error, from the **Error Message** window, select the software and click **Detail**.

Related concepts

- License key types on page 48

Related tasks

- Removing a software application on page 51

Enabling a license

You can enable a license that is in disabled status.

Procedure

1. On the **Resources** tab, click **Storage Systems**, and then expand **All Storage Systems** and the target storage system.
2. Click **License Keys**.
3. In the **License Keys** window, select one or more licenses to enable, and then click **Enable** or **Enable Licenses**.
4. Check the settings in the **Enable Licenses** window, and then click **Apply**.

Related concepts

- License key types on page 48

Related tasks

- Updating license status on page 51
- Disabling a license on page 50

Disabling a license

You can disable a license that is in enabled status.

Procedure

1. On the **Resources** tab, click **Storage Systems**, and then expand **All Storage Systems** and the target storage system.
2. Click **License Keys**.

3. In the **License Keys** window, select one or more licenses to disable, and then click **Disable** or **Disable Licenses**.
4. Check the settings in the **Disable Licenses** window, and then click **Apply**.

Related tasks

- Enabling a license on page 50

Removing a software application

Procedure

1. On the **Resources** tab, click **Storage Systems**, and then expand **All Storage Systems** and the target storage system.
2. Click **License Keys**.
3. In the **License Keys** window, select one or more licenses to uninstall.

> **Note:** On rare occasions, a software option that is listed as **Not Installed** but still has available licensed capacity (shown as ***XX* TB**) might remain in the list. In this case, select that option, and then uninstall the software.

4. Click **Remove** to display the **Remove Licenses** window.
5. Check the settings, and then click **Apply**.

> **Note:** To reinstall a license key after uninstalling it, contact customer support to reissue the license key file.

Related tasks

- Installing the software on page 49

Updating license status

In the following cases, the status of software might remain at Not Enough License or Grace Period. In that case, update the license status.

- When the licensed capacity exceeds the mounted capacity after you reduce the number of LDEVs
- When the licensed capacity exceeds the used capacity after you delete pairs or pool volumes

Procedure

1. On the **Resources** tab, click **Storage Systems**, and then expand **All Storage Systems** and the target storage system.
2. Click **License Keys**.

3. In the **License Keys** window, click **Update License Status**.
4. Check the settings, and then click **Apply**

Related tasks

- Enabling a license on page 50

Viewing license information

You can view license information for software products.

Procedure

1. On the **Resources** tab, click **Storage Systems**, and then expand **All Storage Systems** and the target storage system.
2. Click **License Keys**.

Examples of license information

The following table provides examples of license information displayed in the **License Key** window.

License key status (example)	Status	Key type	Licensed capacity	Term (Days)
Not installed	Not installed	blank	Blank	Blank
Installed with the permanent key	Installed	permanent	Permitted	-
Installed with the term key and set to Enabled	Installed	term	Permitted	Number of remaining days before expiration
Installed with the term key and set to Disabled	Installed (Disabled)	term	Permitted	-
Installed with the temporary key.	Installed	temporary	-	Number of remaining days before expiration
Installed with the emergency key.	Installed	emergency	-	Number of remaining days before expiration
A temporary key was installed, but has expired.	Expired	temporary	-	Number of remaining days before expiration
A term key or an emergency key was installed, but has expired.	Not installed	blank	Blank	Blank
Installed with the permanent key or the term key, but the licensed capacity was insufficient.	Not Enough License	permanent or term	Permitted and Used	-
Installed with the permanent or term key, and then LDEVs are added,	Grace Period	permanent or term	Permitted and Used	Number of remaining days before expiration

License key status (example)	Status	Key type	Licensed capacity	Term (Days)
but the license capacity was insufficient.				
Installed with the temporary key, and then reinstalled with the permanent key, but the license capacity was insufficient.	Installed	temporary	Permitted and Used	Number of remaining days before expiration
Installed with the permanent or term key, then reinstalled with the emergency key.	Installed	emergency	Permitted and Used	Number of remaining days before expiration

Related concepts

- License key types on page 48

Related references

- Using reports to verify system changes on page 500

3

Discovering, registering, and adding management targets

Managing storage resources requires them to be registered in Hitachi Command Suite. Registration automatically occurs after discovery.

- ☐ Setting up storage resources
- ☐ Registering storage systems
- ☐ Registering hosts
- ☐ Registering file servers

Setting up storage resources

Collectively, storage systems, file servers, and hosts are called storage resources. To manage storage resources by using HCS, they must first be discovered and registered.

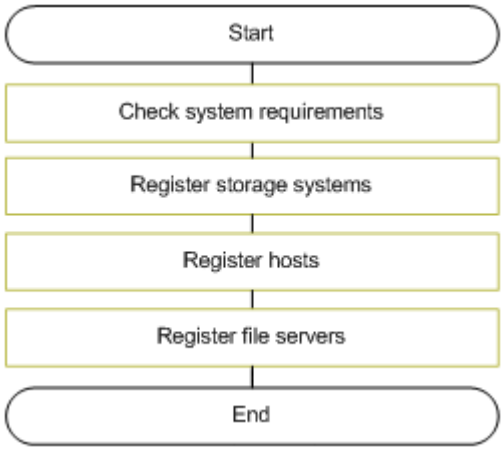
Registering storage resources requires each resource to be reachable on the network by its IP address. After storage and host resources are registered, you can begin allocating volumes to hosts.

Before discovering each resource, verify the system prerequisites and verify setting requirements for the corresponding system configurations described in the *Hitachi Command Suite Administrator Guide*.

Configure the prerequisite environment based on specific storage system requirements. For example, install and enable software licenses by using Storage Navigator, Storage Navigator Modular 2, or other storage management tools, or configure a network based on the system configuration.

If Device Manager is set up in an environment that has already been configured using Storage Navigator, and if you want to use volume labels and DP pool names from a storage system in Device Manager, register the storage system in Device Manager, and then apply the storage system resource labels to Device Manager.

The following figure describes a basic task flow for setting up storage resources.



Registering storage systems

Before you can manage a storage system's resources, you must first register the storage system.

About registering and removing a storage system

Registering a storage system in HCS is done when it is discovered, which then enables you to manage its resources, including:

- Parity Groups
- LDEVs
- Volumes
- DP pools
- Resources available through external storage connections

When you specify an IP address (or host name) and authentication information for a storage system, the storage system is discovered and registered based on the specified information. For Storage Management Initiative Specification (SMI-S) enabled storage systems, the SMI-S provider's IP address is used.

You can modify information about a storage system after a storage system is registered.

If managing a storage system with HCS is no longer necessary (for example because of reorganization), you can remove the registered storage system. Removing a storage system results in the following:

- Configuration information such as parity groups, volumes, and allocated paths are retained within the storage systems.
- Any information configured or specified in HCS, such as tiers and data placement profiles, is deleted.

Related concepts

- About acquiring the most recent storage system information on page 63

Related tasks

- Registering a storage system on page 60
- Changing storage system information on page 62
- Removing a storage system on page 63

Related references

- Prerequisites for registering a storage system on page 57

Prerequisites for registering a storage system

Before you can discover and register a storage system, collect storage system-related information and complete the following prerequisite tasks.

Collect and confirm the following storage system information:

Storage system type	Required information
Universal Storage Platform V/VM	• Storage system type • Storage system IP address (IPv4 or IPv6) or host name (alias name)
Virtual Storage Platform	• Storage system type • Storage system IP address (IPv4 or IPv6) or host name (alias name)*
VSP G1000, VSP G1500, VSP F1500	• Storage system type • Storage system IP address (IPv4 or IPv6) or host name (alias name)*
VSP Gx00 models, VSP Fx00 models	• Storage system type • Storage system IP address (IPv4 or IPv6) or host name (alias name)* • Port number: The RMI registry port number of the SVP to use for communication from HCS to SVP (when changing to a value other than the default value) • URL of NAS Manager (when using Virtual Storage Platform G400, G600, G800, and Virtual Storage Platform F400, F600, F800 with NAS modules)
HUS VM	• Storage system type • Storage system IP address (IPv4 or IPv6) or host name (alias name)*
* For VSP Gx00 models, VSP Fx00 models, VSP G1000, VSP G1500, VSP F1500, Virtual Storage Platform, or HUS VM, if SSL/TLS communication is used between the storage system or SVP and management client, specify the host name set in Common Name in the server certificate to register the storage system. For details about secure communications, see the *Hitachi Command Suite Administrator Guide*.	

After you collect the required information, complete the following required tasks:

Storage system type	Required tasks
Universal Storage Platform V/VM	• Connect the target storage system to the network. • Configure the Fibre Channel and iSCSI ports (including the topology and data transfer speed). • Create accounts in Device Manager - Storage Navigator.[1]
Virtual Storage Platform	• Connect the target storage system to the network. • Configure the Fibre Channel and iSCSI ports (including the topology and data transfer speed). • Create accounts in Device Manager - Storage Navigator.[1] • Configure the SSL/TLS communication settings (if they are to be configured

Storage system type	Required tasks
	between the management client and storage system).[2]
VSP G1000, VSP G1500, VSP F1500	• Connect the target storage system to the network. • Configure the Fibre Channel and iSCSI ports (including the topology and data transfer speed). • Create accounts in Device Manager - Storage Navigator.[1] • Configure the SSL/TLS communication settings (if they are to be configured between the management client and storage system).[2] • Configure the SSL/TLS communication settings (if they are to be configured between the HCS management server and storage system).[3]
VSP Gx00 models, VSP Fx00 models	• Connect the target storage system to the network. • Configure the Fibre Channel and iSCSI ports (including the topology and data transfer speed). • Create accounts in Device Manager - Storage Navigator.[1] • Configure the SSL/TLS communication settings (if they are to be configured between the management client and storage system).[2] • Configure the SSL/TLS communication settings (if they are to be configured between the HCS management server and storage system).[3] • Startup of the RMI server on the SVP (when multiple storage systems are registered for a single SVP).
HUS VM	• Connect the target storage system to the network. • Configure the Fibre Channel and iSCSI ports (including the topology and data transfer speed). • Create accounts in Device Manager - Storage Navigator.[1] • Configure the SSL/TLS communication settings (if they are to be configured between the management client and storage system).[2]

1. For VSP G1000, VSP G1500, VSP F1500, Virtual Storage Platform, HUS VM, VSP Fx00 models or VSP Gx00 models, you must select the Administrator user group (which is a built-in group) to register the accounts. For VSP Fx00 models or VSP G400, G600, G800 with NAS modules (of microcode version earlier than 83-04-0X-XX/XX), you must also select the Support Personnel user group.

2. The name of the storage system must be resolved from the management client.

3. When you use HCS to manage software such as Device Manager - Storage Navigator or RAID Manager and users of this software, SSL/TLS communication must be configured between the

Storage system type	Required tasks
	HCS management server and the storage system or SVP. Note that, for VSP Gx00 models or VSP Fx00 models, HCS always manages users and for VSP G1000, VSP G1500, or VSP F1500, you can choose whether HCS will manage users. If HCS is to manage users, specify the following items in advance: -Register a firewall exception (check the port to be used). -Resolve the name of the HCS management server from the storage system (SVP). Even if you do not configure the settings for SSL/TLS communication, when HCS is installed, the default certificate enables the settings for encrypted communication. For details about ports to be used and how to configure the settings, see the *Hitachi Command Suite Administrator Guide*.

Related concepts

- About registering and removing a storage system on page 57
- About user accounts and controlling access to resources on page 88

Related tasks

- Registering a storage system on page 60
- Setting the Fibre Channel port address on page 237
- Setting the data transfer speed on a Fibre Channel port on page 235

Registering a storage system

A storage system must be registered in Hitachi Command Suite before it can be managed. After a storage system is discovered and registered, it becomes a managed resource.

Before you begin

Review the prerequisites for your storage system type before performing this task.

Procedure

1. On the **Administration** tab, select **Managed Resources**.
2. On the **Storage Systems** tab, click **Add Storage System**.
3. Specify the storage system type and then provide the required information, depending on the storage system type that you choose:

Storage system type	Required settings
All storage system types	• An IP address (or hostname). • Authentication information (username and password) for most supported storage systems.
VSP Fx00 models, VSP Gx00 models	• IP address (IPv4 or IPv6) or host name (alias name) of the SVP. • Port number: The RMI registry port number of the SVP to use for communication from HCS to SVP (when changing to a value other than the default value).

Storage system type	Required settings
	• User name and password for the storage system.
VSP G1000, VSP G1500, VSP F1500	Enable user account authentication. If authentication of user accounts is enabled: • User accounts are authenticated by HCS when they log in to CCI and SVP. • User accounts that are created in HCS can be used in CCI and in SVP, allowing centralized management of user accounts. • From CCI or the SVP you will be able to perform tasks (such as, adding users to a user group, and assigning resource groups and roles to a user group) according to the access control settings made in HCS. If you want to create a new HCS user account for performing tasks on CCI or the SVP, in HCS set up access control for the storage resources. **Note:** In the event of the unavailability of the HCS server, CCI authentication will be forwarded directly to the SVP. If you register the same user accounts for HCS into the SVP, you can use the user account for the SVP to perform CCI authentication even if the HCS server is not running.
For mid-range storage systems	• Controller (0,1) • Protocol information (secure, not secure)
For SMI-S enabled storage systems	• SMI-S provider IP address • Authentication information • Protocol information (secure, not secure) • And, possibly, specify a non-default port number If an SMI-S provider that manages multiple SMI-S enabled storage systems is specified, all the SMI-S enabled storage systems under that provider are registered.

4. For VSP Fx00 models and VSP Gx00 models, specify the required settings.
5. Click **OK**.
6. Check the **Data Collection Tasks** tab to verify that the task completes successfully.
7. When the registration task completes, confirm that the registered storage system appears on the **Storage Systems** tab in the **Managed Resources** list, or in the **Storage Systems** tree on the **Resources** tab.
8. If you are using VSP G400, G600, G800, and VSP F400, F600, F800 with NAS modules, register Device Manager information by using NAS Manager.
   Device Manager users that are specified in NAS Manager must belong to the **PeerGroup** as a user group. As a best practice, do not use the default **HaUser** user account. Instead, create a new account in Device Manager.
9. Use NAS Manager to synchronize the information between the NAS module and Device Manager.

Result

The storage system is now a managed resource.

Related concepts

- About registering and removing a storage system on page 57
- About user accounts and controlling access to resources on page 88

Related tasks

- Changing storage system information on page 62
- Removing a storage system on page 63

Related references

- Prerequisites for registering a storage system on page 57

Changing storage system information

You can modify the IP address, host name, or login information (user ID and password) for a storage system. You can also modify the SVP IP address or enable user account authentication for VSP Fx00 models and VSP Gx00 models. You can select multiple storage systems to modify at one time, but when selecting multiple storage systems, only login information (user ID and password) can be modified.

When you change the user account or password for a storage system, you must first make the changes in Hitachi Device Manager - Storage Navigator (or Storage Navigator), and then make those same changes from the Edit Storage Systems dialog box using HCS.

For more information about changing the user account of storage systems, see the *System Administrator Guide* or the *Storage Navigator User Guide* for your storage system.

Procedure

1. On the **Administration** tab, select **Managed Resources**.
2. On the **Storage Systems** tab, select the storage system, or storage systems, you want to modify.
3. Click **Edit Storage Systems**.
4. Specify the appropriate information and click **OK**.
5. Check the **Data Collection Tasks** tab to verify that the task has completed successfully.
6. Confirm that the information in the storage system list is updated.

Tip: For VSP Fx00 models and VSP Gx00 models, refresh the storage system to update the latest information after you change the IP address of the storage system controller on the SVP.

Note: For VSP G400, G600, G800 and VSP Fx00 models with NAS modules: If you use an account that was created in a firmware version earlier than 83-03-2X-XX/XX and the firmware was updated to version 83-03-2X-XX/XX or later, then change the password. Otherwise, you will be unable to operate the NAS module.

Related concepts

- About registering and removing a storage system on page 57

Related tasks

- Registering a storage system on page 60

Removing a storage system

If managing a storage system with Hitachi Command Suite is no longer necessary (for example because of reorganization), you can remove the registered storage system.

Procedure

1. On the **Administration** tab, select **Managed Resources**.
2. On the **Storage Systems** tab, select one or more storage systems.
3. Click **Remove Storage Systems**.
4. When the task completes, confirm that the removed storage system does not appear on the **Storage Systems** tab of **Managed Resources**, or in the **Storage Systems** tree on the **Resources** tab.

Result

The storage system is removed.

Related concepts

- About registering and removing a storage system on page 57

About acquiring the most recent storage system information

To get the most current information about a target storage system, manually refresh the system to display it.

When you click Refresh Storage System to refresh it, the following information is updated for storage resources:

- Volumes
- Parity groups
- DP Pools
- External storage connections

Tip: If you mapped volumes in the external storage system by using CCI or Storage Navigator, you can confirm the mapped volumes from the internal storage system by refreshing storage system information in the following order:

- External storage system
- Internal storage system

When data placement profiles for HDT volumes are being managed, you can specify an option in the Refresh Storage Systems dialog box to search for inconsistencies in data placement profiles.

To update information about hosts, such as the host name, WWN, or iSCSI name, refresh the hosts.

To update host bus adapter (HBA) WWN nicknames, refresh one or more storage systems. Refreshing host information does not update WWN nicknames. When several WWN nicknames are assigned to a single HBA, only one of the nicknames is displayed for that HBA.

If connected to Hitachi Tuning Manager, users can also update performance information that is displayed in the Mobility tab. To automatically update the storage system information, specify proper settings in the `server.properties` file or `dispatcher.properties` file of the management server. For details about specifiable values, see the *Hitachi Command Suite Administrator Guide*.

Related concepts

- About registering and removing a storage system on page 57

Related tasks

- Acquiring the most recent storage system information on page 64
- Viewing current storage system information on page 523

Acquiring the most recent storage system information

Refreshing a storage system updates and displays the latest storage system information.

Before you begin

For VSP Fx00 models and VSP Gx00 models, if you change the IP address of the storage system controller, you must update the storage system information on the SVP before you refresh the storage system.

Procedure

1. On the **Administration** tab, select **Managed Resources**.
2. On the **Storage Systems** tab, select one or more storage systems to refresh.

3. Click **Refresh Storage Systems**, then click **OK**.
4. Check the **Data Collection Tasks** tab to verify that the task completed.

Result

Updated storage system information displays in the list of storage systems. Updated storage system information is also displayed on the Resources tab, in the Storage Systems tree.

Note: Label information for volumes and DP pools is not refreshed. To apply existing label information from a storage system to HCS, click Refresh Labels. For VSP G1000, VSP G1500, VSP F1500, VSP Fx00 models and VSP Gx00 models, refreshing labels is unnecessary since label information from the storage system to HCS will always match.

Related concepts

- About acquiring the most recent storage system information on page 63

Operations available to SMI-S enabled storage systems

Discovering and registering storage systems that support an SMI-S provider for storage system management enables you to perform certain tasks.

The Storage Management Initiative Specification (SMI-S) defines software development standards in support of storage management interoperability. An SMI-S provider is the software component that resides on a resource, such as a storage system, providing network management (over HTTP/HTTPS) by management clients, such as Hitachi Command Suite.

When you manage storage systems that support an SMI-S provider, the SMI-S enabled storage systems become Hitachi Command Suite resources on which you can perform the following tasks:

- Register, change, delete, and update the SMI-S enabled storage systems.
- Display summary information for SMI-S enabled storage systems.
- Operate SMI-S enabled storage systems from the SMI-S enabled storage system management tool (launch the management tool from Element Manager).
- Externally connect and virtualize an SMI-S enabled storage system for migrating data.

Tip: Virtualizing an SMI-S enabled storage system without registering the system in HCS can also be done by using the port scan and volume discovery functionality.

You must be assigned All Resources to view or perform tasks on SMI-S enabled storage systems.

Related tasks

- Discovering and virtualizing volumes of an unregistered storage system on page 163

Registering hosts

You can register hosts as managed resources using Device Manager so that storage system volumes can be used by hosts and their applications.

About registering a host

Discover and register hosts to manage them in Hitachi Command Suite.

The following host types can be added to HCS:

- Normal hosts (an open host that is not virtualized)
- Virtualization servers (a physical server on which virtualization software is installed)
- Virtual machines (a machine that is created on the virtualization server and uses virtualization software)

Depending on the purpose, there are multiple methods for discovering and adding hosts, including manually, by using Host Data Collector (agentless discovery), or by using a Device Manager agent. For host prerequisites and environment settings, see the *Hitachi Command Suite Administrator Guide*.

You can also monitor the usage status of storage resources from the host point of view by generating reports that reflect the actual capacity being used by each host.

Related concepts

- Methods for registering hosts on page 66
- About changing host settings and information on page 73

Related references

- Priority for acquiring the WWN or iSCSI name on page 68

Methods for registering hosts

There are a variety of methods to register one or more hosts including using its WWN or iSCSI name, using host-installed agents, or agentless discovery.

The following table describes the methods you can use to register one or more hosts:

Table 1 Methods for registering hosts

Purpose	Method	Description
To manually register individual hosts.	Register each host by WWN/ iSCSI name.	To register hosts in HCS based on WWN/iSCSI name, manually specify the host name, OS type, and WWN/iSCSI name information. Hosts can be registered without Host Data Collector or Device Manager agent setup. Virtualization servers can also be manually registered, but the virtual machine information cannot be referenced or displayed.
To register multiple hosts in a batch.	Use Host Data Collector.	Register multiple hosts in batch by specifying the IP address (range specification and listing addresses are available), and the authentication information of the host on the network. The host's WWN/iSCSI name information is automatically obtained. To allocate volumes to an FCoE port, you must manually add a WWN. Host Data Collector setup is required.
To know the usage status of storage resources for the virtualization servers, or the correspondence with the virtual machines.	Use Host Data Collector.	To know the usage status of storage resources for the virtualization servers, or the correspondence with the virtual machines, and when you want to register both the virtualization server and the virtual machine, the virtualization server must use NPIV (N-Port ID Virtualization) HBAs, and be connected with Fibre Channel or FCoE. If the virtualization server is not using NPIV HBAs, either a virtualization server or a virtual machine can be registered per HBA unit (both a virtualization server and a virtual machine cannot be registered in the same WWN).
To newly install Hitachi Command Suite to an existing environment that was configured with Hitachi Storage Navigator.	Scan a host.	Use Host Scan when you want to newly install HCS to an existing environment that was configured by using Hitachi Storage Navigator.
To synchronize host information between Device Manager and Compute Systems Manager.	Install Device Manager and Compute Systems Manager on the same server, and obtain host information through Host Data Collector.	If Device Manager and Compute Systems Manager are installed on the same server, information about hosts being managed by both products is synchronized automatically when using Host Data Collector. Therefore, hosts registered in Compute Systems Manager are also registered in Device Manager. However, host deletion is not synchronized. If you want to re-synchronize a Compute Systems Manager host that has been deleted from Device Manager, you must manually re-synchronize the managed hosts.

As an alternative method for registering hosts to HCS, install the Device Manager agent on each host or on the pair management server. Information, such as host information and copy pair information, will be sent from the

Device Manager agent to the HCS server. However, if you want to allocate volumes to an FCoE port, you must manually add a WWN.

To exclude a host from being managed by Device Manager, uninstall the Device Manager agent from the host, and then delete the host from Device Manager.

Tip: To register multiple hosts in a batch operation, use the `ImportHosts` command to import host information such as host name and WWN/iSCSI name, from a CSV file, and then register multiple hosts in a batch operation. Hosts that are registered by using the `ImportHosts` command are managed as manually registered hosts. For details about the `ImportHosts` command, see the *Hitachi Command Suite CLI Reference Guide*.

Related concepts

- About registering a host on page 66
- About linking with Compute Systems Manager on page 577

Related tasks

- Registering hosts by using Host Data Collector on page 69
- Registering hosts manually by specifying the WWN or iSCSI name on page 70
- Registering hosts using host scan on page 71

Related references

- Priority for acquiring the WWN or iSCSI name on page 68

Priority for acquiring the WWN or iSCSI name

If hosts you want to register include a host that has a WWN or an iSCSI name that has already been registered, the host that acquires the WWN or iSCSI name is determined based on the product or application that detects it, and is registered according to the following acquisition priority list.

When hosts are registered by using multiple methods, the WWN or iSCSI name acquisition priority, listed from highest to lowest is:

1. A virtualization server (other than one registered manually) (iSCSI is not supported).
2. A virtualization server that is registered manually (iSCSI is not supported for virtualization servers).
3. A host registered by using Device Manager agent.
4. A normal host or a virtual machine registered either manually or by using Host Data Collector.

When hosts with the same WWN or iSCSI name are registered by using the same method:

- If a host is registered without using the Device Manager agent, the WWN or iSCSI name of the host that was registered first is retained.

- If a host is registered by using the Device Manager agent, the WWN or iSCSI name is registered to the host that was registered last.

Related concepts

- About registering a host on page 66
- Methods for registering hosts on page 66

Registering hosts by using Host Data Collector

Use Host Data Collector to discover and register multiple hosts, virtualization servers, and virtual machines in a single task.

When using Host Data Collector, if Device Manager is linked with Compute Systems Manager, hosts managed by Compute Systems Manager will be registered in Device Manager.

Before you begin

- Gather the following information:
  - Host OS type
  - Host IP address (IPv4)

    Note: To detect multiple hosts at the same time, specify a subnet or a range of 256 or fewer IP addresses, or specify individual IP addresses.

  - User name and password for logging on to the host (with Administrator or superuser permissions)

    Note: When specifying users other than the root, you need to specify environment settings for when general users are used to collect host information on the host side. For details about how to specify settings, see the *Hitachi Command Suite Administrator Guide*.

  - IP address of Host Data Collector (IPv4 or IPv6) or Host Name
- Gather the following information when registering a virtualization server:
  - IP address of Host Data Collector (IPv4 or IPv6) or Host Name
  - IP address of VMware vCenter Server or a virtualization server (IPv4 or IPv6)

    Note: To register a virtualization server, you cannot specify a subnet, a range, and multiple IP addresses.

  - User name and password of a user (who has the role of system administrator) for logging on to vCenter Server or virtualization servers

Note: If vCenter Server is being used, enter the user name and password for vCenter Server. If vCenter Server is not being used, you need to enter and register user names and passwords of virtualization servers one by one if multiple virtualization servers are being managed.

- Verify that target hosts are connected to the network.
- Specify environment settings for virtualization servers and virtual machines. For detailed prerequisites see the *Hitachi Command Suite Administrator Guide*.

Procedure

1. On the **Administration** tab, select **Managed Resources**.
2. On the **Hosts** tab, click **Add Hosts**.
3. In the **Add Hosts** dialog box, click **Discover Hosts**.
   If you are registering a virtualization server, select **VMware** from **OS Type**.
4. Specify the necessary items, select the **Host Data Collector** to be used, and then submit the task.
   The task is registered and listed on the **Data Collection Tasks** tab.
5. Verify task status on the **Data Collection Tasks** tab.

Result

Registered hosts can be viewed in the Resources list on the Navigation pane.

Tip: If you want to allocate volumes to an FCoE port, you need to manually add a WWN.

Related concepts

- Methods for registering hosts on page 66

Registering hosts manually by specifying the WWN or iSCSI name

You can register hosts manually to manage and monitor the hosts by specifying a WWN or iSCSI name.

Before you begin

- Host name
- OS type
- Unregistered WWN or iSCSI name

Procedure

1. On the **Administration** tab, select **Managed Resources**.

2. On the **Hosts** tab, click **Add Hosts**.
3. In the **Add Hosts** dialog box, select **Add Hosts Manually**.
4. Specify the appropriate information, and then click **OK** to submit the task.

   If specifying a virtualization server, in **OS Type** select **VMware**.

Result

All registered hosts, virtual machines, and virtualization servers are displayed by each host OS type (for example: AIX, HP-UX, Linux, Solaris, Windows, Virtualization Servers, or Others) in the Resources list on the Navigation pane.

Related concepts

- Methods for registering hosts on page 66

Registering hosts using host scan

You can register hosts by using host scan to discover hosts so that they can be managed and monitored.

Scanning for hosts automatically creates hosts with host names that match the WWNs or iSCSI names that are registered in the host group or iSCSI targets on the storage system. For best results, your environment should satisfy the following conditions:

- The name of a host group or iSCSI target should be the same as the actual host, and the WWN or iSCSI name for that host is set for LUN security.
- Only a single host can be related to each host group or iSCSI target.

Before you begin

- Install LUN Manager on the target storage system.

Procedure

1. On the **Administration** tab, select **Managed Resources**.
2. On the **Hosts** tab, click **Host Scan**.
3. Confirm that your environment meets the recommended conditions, and then select storage systems on which to scan for hosts.
4. Click **Submit**.

Result

All registered Data Collection Tasks hosts, virtual machines, and virtualization servers can be viewed by host OS type (for example: AIX, HP-UX, Linux, Solaris, Windows, Virtualization Servers, or Others) in the Resources list on the Navigation pane. If a virtualization server that has an iSCSI name is scanned, it is displayed in the host OS type, Others.

Tip: If one host has multiple WWNs or iSCSI names detected, even if it is one host in the actual environment, it might occasionally be registered as multiple hosts by Device Manager. This is because hosts are generated in host groups or iSCSI targets when one host has multiple WWNs or iSCSI names. In this case, the hosts must be merged, meaning merge the WWNs or iSCSI name. Merge the hosts that have been added as multiple hosts in Device Manager into one host, meaning that WWN or iSCSI name information is merged, and the source hosts are deleted automatically.

Related concepts

- Methods for registering hosts on page 66

Related tasks

- Registering hosts using merge hosts on page 72
- Installing the software on page 49

Related references

- Workflow for detecting hosts on page 74

Registering hosts using merge hosts

You can use merge hosts to combine source and target host information into a single surviving target host that is registered. The maximum number of WWNs and iSCSI names that you can merge at one time is 100.

Before you begin

- Identify the names of the source and target host. Note that the source host is merged into the target host, then the source host is deleted.
- Ensure the source host has been registered either manually or by host scan.
- Ensure the target host is not a mainframe host, a file server, or a virtualization server (except for manually registered virtualization servers).
- Ensure the target host has at least one registered WWN or iSCSI name.

Procedure

1. On the **Administration** tab, select **Managed Resources**.
2. Select the **Hosts** tab, select the target host row, click **More Actions** and select **Merge Hosts**.
3. In the merge hosts window, select the source host, then click **Show Plan**.
4. Review the target (surviving host) and source (host to be deleted) information.
   - If the target and source host information is correct, click **Submit**.
   - If the target or source host information must be changed, click **Back**.
   - To abandon the task, click **Cancel**.

5. On the **Resources** tab, in the hosts tree, confirm the target host information and the source host removal. You can also confirm this on the **Administration** tab in the **Hosts** tree under **Managed Resources**.

Related concepts

- Methods for registering hosts on page 66

About changing host settings and information

You can change host settings and update previously registered host information.

Depending on how a host is registered, you change host settings and information by using the Refresh Hosts or Edit Hosts options. The Refresh Hosts option on the Hosts tab updates a host by obtaining information from the Host Data Collector or Device Manager agent. The Edit Hosts option allows you to manually change host settings.

On the vCenters tab, Refresh vCenters can be used to refresh information for virtualization servers that are managed by VMware vCenter Server.

Note: If a Device Manager host is registered on Compute Systems Manager, and if you update host information using Compute Systems Manager, that information is automatically updated in Device Manager.

You can update other host-related information when replacing host HBAs by exchanging the WWN settings in the LUN paths so that the new HBA inherits the WWN settings from the old HBA. Perform this settings exchange by using the Exchange HBAs dialog box when you replace a host HBA due to failure, or by using the Add HBAs dialog box and the Remove HBAs dialog box when you complete a planned HBA replacement as part of scheduled maintenance to prevent performance degradation.

Tip: The settings for the host WWN or iSCSI name are not linked to the settings for the host group or iSCSI target on the storage system side. For example, even if you delete the host WWN or iSCSI name, the WWN or iSCSI name is not deleted from the host group or iSCSI target.

To delete the WWN after replacing HBAs, use the Remove HBAs dialog box. To delete the WWN or iSCSI name from a host group or an iSCSI target at a time other than when replacing an HBA, use the `ModifyPort` command.

For details about using the `ModifyPort` command, see the *Hitachi Command Suite CLI Reference Guide*.

Related concepts

- About registering a host on page 66

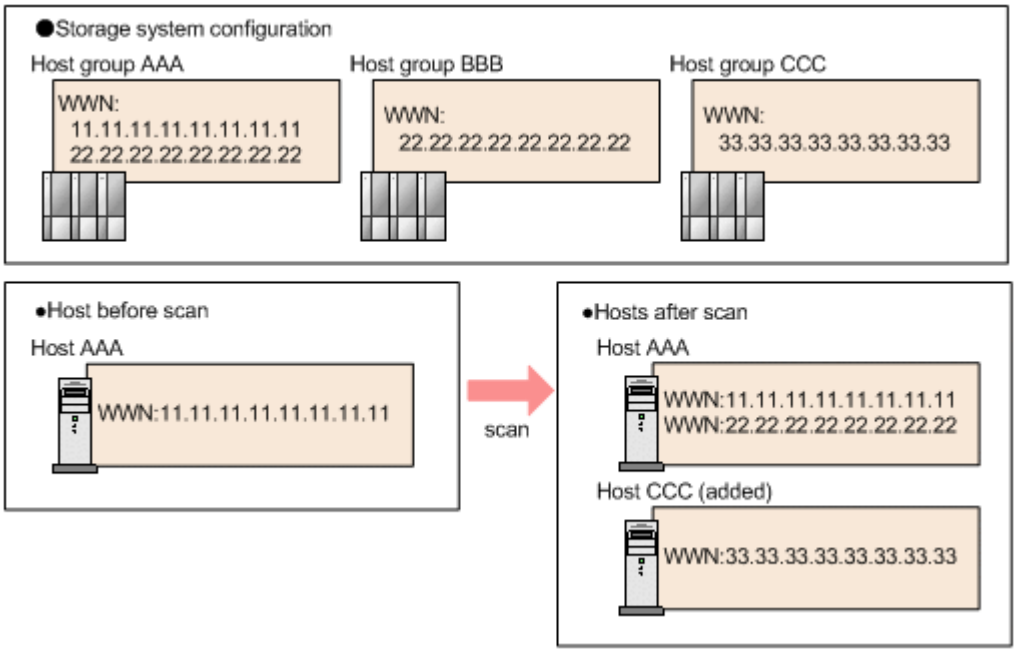
Related tasks

- Updating host information registered by using Host Data Collector on page 75
- Changing settings for a manually registered host on page 76
- Changing settings for a host registered by using Device Manager agent on page 77
- Refreshing vCenter hosts on page 78

Workflow for detecting hosts

This topic describes how hosts are detected and registered in Device Manager. Read this topic if you are installing Hitachi Command Suite in an environment that was configured by using Storage Navigator.

A host group or iSCSI target is configured on the storage system side, and the storage system is registered in Device Manager. After that, if a host is detected by Device Manager, a host group in the storage system that is selected at the time of detection of the host or a host that has the same name as the name of the iSCSI target is registered in Device Manager. If the same WWN is registered to multiple hosts, or if the same iSCSI name is registered to multiple iSCSI targets, only the WWN or iSCSI name with the same name as the host that is already registered in Device Manager is added to the host. The following shows an example of host detection:



In this diagram, the WWN of 22.22.22.22.22.22.22.22 for host group AAA is added to host AAA, which has the same name as the host group. The WWN of host group BBB is already registered to host AAA, and host BBB cannot be registered in Device Manager. Host CCC, which has the same name as host group CCC, is added to Device Manager, and the WWN of 33.33.33.33.33.33.33.33 for host group CCC is registered to host CCC.

If one host that has multiple WWNs or iSCSI names is detected, even if it is one host in the actual environment, it might occasionally be registered as multiple hosts by Device Manager. This is because hosts are generated in host groups or iSCSI targets when one host has multiple WWNs or iSCSI names. In this case, the hosts must be merged.

Merge the WWNs and iSCSI names. Merge the hosts that have been added as multiple hosts in Device Manager into one host. WWN or iSCSI name information is merged, and the source hosts are deleted automatically.

Tip: If one host that has multiple WWNs or iSCSI names is detected, even if it is one host in the actual environment, it might occasionally be registered as multiple hosts by Device Manager. This is because hosts are generated in host groups or iSCSI targets when one host has multiple WWNs or iSCSI names. In this case, the hosts must be merged. Merge the WWNs or iSCSI names. Merge the hosts that have been added as multiple hosts in Device Manager into one host. WWN or iSCSI name information is merged, and the source hosts are deleted automatically.

Related tasks

- Registering hosts using host scan on page 71

Updating host information registered by using Host Data Collector

You can update information for registered hosts using Host Data Collector. If Device Manager is linked to Compute Systems Manager, updated host information is synchronized between Device Manager and Compute Systems Manager.

Use Refresh Hosts to update host information registered by using Host Data Collector.

To change authentication information about a registered IP address or user account, edit the IP address or user account by using the Edit Hosts dialog box, and then refresh host information.

To make changes to each item, use the method in the following table:

Table 2 Updating host information using Host Data Collector

Item	Refresh Hosts	Edit Hosts
Host Name	Y	N
OS Type	N	N
Host IP Address	N	Y[1]
User ID/PW	N	Y
Host Data Collector IP Address	N	Y[2]
Protocol	N	Y

Item	Refresh Hosts	Edit Hosts
Port Type	Y[3]	Y[3,5]
WWN/iSCSI name	Y[3]	Y[3,4,5]
WWN Nickname	N[6]	N

Legend:
- Y : Can be edited or refreshed
- N : Cannot be edited or refreshed

Notes:
1. For virtualization servers, you can specify either IPv4 or IPv6.
2. You can specify either IPv4 or IPv6.
3. Clicking Refresh Hosts will discover new or updated WWNs/iSCSI names and their port types. Deleted WWNs/iSCSI names require that you click Edit Hosts and delete the WWNs/iSCSI names that need to be removed.
4. If you want to allocate volumes to a FCoE port, you need to manually add a WWN.
5. You cannot specify this for a virtualization server.
6. To update WWN nicknames that have been specified by using storage system management tools such as Storage Navigator, refresh the storage system information. When several WWN nicknames are assigned to a single HBA, only one of the nicknames is displayed for that HBA.

Procedure

1. On the **Administration** tab, select **Managed Resources**.
2. On the **Hosts** tab, select the host to update, then click either **Refresh Hosts**, or **Edit Hosts**, depending on the items that you want to update.
3. If necessary, specify updated information and click **OK**. The task is registered and added to the task list on the **Data Collection Tasks** tab. If you updated information in the **Edit Hosts** dialog box, after the task has completed successfully, refresh the hosts using **Refresh Hosts**.



Tip: On the **vCenters** tab, **Refresh vCenters** can be used to refresh information for virtualization servers that are managed by VMware vCenter Server.

Result

The host list is updated.

Related concepts

- About changing host settings and information on page 73

Related tasks

- Refreshing vCenter hosts on page 78

Changing settings for a manually registered host

You can use Edit Hosts to update host information registered by specifying WWNs/iSCSI names, or registered by using the host detection function.

Procedure

1. On the **Administration** tab, select **Managed Resources**.
2. On the **Hosts** tab, select the host to change, and click **Edit Hosts**.
3. Specify the required items, and then click **OK**.

Result

The host information is updated.

Related concepts

- About changing host settings and information on page 73

Changing settings for a host registered by using Device Manager agent

When the Device Manager agent is used, host information is periodically sent from the Device Manager agent. To update displayed host information, use the Refresh Hosts button. Additionally, after changing settings on the Edit Hosts dialog box, depending on which settings you changed, it may be necessary to use Refresh Hosts to update host information.

To make changes or update host information to each item, see the following table:

Table 3 Updating host information using Device Manager agent

Item	Refresh Hosts	Edit Hosts
Host Name	N	Y[1]
OS Type	Y	N
IP Address	Y	N
Port Type	Y[2]	Y[2]
WWN/iSCSI name	Y[2]	Y[2,3]
WWN Nickname	N[4]	N

Legend:
- Y : Can be edited or refreshed
- N : Cannot be edited or refreshed

Notes:
1. To change the host name:
   a. On the host, change the host name.
   b. (Optional) Restart the host.
   c. Click Edit Hosts, then change the host name to the same name that you used at step a.
   d. Restart the Device Manager agent service. Wait until the host has finished starting up before restarting the Device Manager agent service.
   If both the old and new host name are displayed, then delete the old host.
2. WWNs/iSCSI names and their port types that have been added can be affected by the Refresh Hosts button, but deleted WWNs/iSCSI names and their port types cannot be affected. To delete such WWNs/iSCSI names, click Edit Hosts.
3. If you want to allocate volumes to a FCoE port, you need to manually add a WWN.

Item	Refresh Hosts	Edit Hosts
4. To update WWN nicknames that have been specified by using other storage system management tools, such as Storage Navigator, refresh the storage system information. When several WWN nicknames are assigned to a single HBA, only one of the nicknames is displayed for that HBA.		

Procedure

1. From the **Administration** tab, select **Managed Resources**.
2. On the **Hosts** tab, select the host to update, then click either **Refresh Hosts** or **Edit Hosts** depending on the items that you want to update.
3. Modify the settings and click **OK**.
   The task is registered to **Data Collection Tasks** tab. If you changed settings in the **Edit Hosts** dialog box, after the task successfully completes, refresh the hosts using **Refresh Hosts**.

Result

The host list is updated.

Note: After you change a host name, both host names (before and after the change) might display on the Resources tab. In this case, delete the host before making the change. When copy pairs are managed by the host where Device Manager agent is installed, in addition to deleting the host, the storage system also needs to be updated.

Related concepts

- About changing host settings and information on page 73

Refreshing vCenter hosts

You can refresh information for multiple virtualization servers that are managed by VMware vCenter Server, without the need to authenticate.

Procedure

1. On the **Administration** tab, select **Managed Resources**.
2. On the **vCenter** tab, select the one or more vCenters, and click **Refresh vCenters** to refresh all virtualization servers (ESX hosts) for each vCenter.
3. In the **Refresh vCenters** dialog box, confirm the refresh by clicking **OK**.
4. The task is registered and added to the task list on the **Data Collection Tasks** tab.

Result

The selected vCenters are refreshed.

Related concepts

- About removing hosts and related resources, or vCenters on page 79

About removing hosts and related resources, or vCenters

For VMware vCenter Servers:

- If all virtualization servers managed by VMware vCenter Server are deleted, the VMware vCenter Server is automatically deleted from the list of VMware vCenter Servers in the vCenters tab.

For hosts:

When hosts no longer need to be managed, for example due to a reorganization, you can remove the registered hosts and exclude them from HCS management. Related storage resources, such as volumes allocated to the target host and path setting information, can also be removed. However, storage resources that are allocated to other hosts are not released. When removing storage resources, you can select to:

- Unallocate volumes
- If the volumes that are unallocated when a host is removed are global-active device paired volumes, the following information is released at the same time:
  - Copy pairs
  - Copy groups
  - Configuration definition files (HORCM)
- Delete associated host groups or iSCSI targets
- Release LUSE volumes
- Delete virtual LDEV ID. This applies when managing resources that have virtual IDs set.
- Delete associated volumes
- Shred associated volumes. Format the volumes if shredding is unavailable.

By automatically executing these options as an extension of removing a host, the process is simplified.

Note: When removing storage resources related to a host, verify that there is not any related data remaining on the volumes.

Tip: If you are removing a host that was registered using Device Manager agent, uninstall the Device Manager agent from the host, then remove the host using HCS.

Note: If host information is synchronized between HCS and Compute Systems Manager, the removal of hosts by HCS is not reflected in Compute Systems Manager.

Related concepts

- About unallocating volumes on page 226
- About managing storage resources that have virtual IDs on page 566
- About deleting unallocated volumes on page 156
- About releasing a LUSE volume on page 159

Related tasks

- Removing hosts and releasing associated resources on page 80
- Refreshing vCenter hosts on page 78
- Unallocating global-active device pairs on page 360
- Unallocating individual global-active device volumes on page 362

Removing hosts and releasing associated resources

When you no longer need to manage a host, remove the host and release its associated related storage resources, such as allocated volumes and path settings.

Before you begin

- Gather the following information:
  - Target host names
  - Target volumes
- Back up the target volumes
- Stop input and output to the target volumes

Procedure

1. On the **Administration** tab, select **Managed Resources** and click the **Hosts** tab.
2. From the lists of hosts, select the host you want to remove and click **More Actions > Remove Hosts**.
3. In the **Remove Hosts** dialog box, select one or more options from the following:
   - Remove only selected hosts
   - Unallocate volumes and remove hosts

   When you unallocate the volumes that are used by the selected hosts and clean up the volumes, select from the following:
   - Delete associated host groups or iSCSI targets
   - Delete associated volumes
   - Shred associated volumes

   Note: Shred during off hours, such as overnight, so that the shredding process does not adversely affect system

performance. To verify the standard required times for shredding, see the *Hitachi Volume Shredder User Guide*.

- Release LUSE volumes
- Delete virtual LDEV ID. This applies when managing resources that have virtual IDs set.

4. Click **Plan Details** to confirm that the information about associated resources is correct.
5. (Optional) Update the task name and provide a description.
6. (Optional) Expand **Schedule** to specify the task schedule.

   You can schedule the task to run immediately or later. The default setting is **Now**. If the task is scheduled to run immediately, you can select **View task status** to monitor the task after it is submitted.
7. Click **Submit.**
   If the task is scheduled to run immediately, the process begins.
8. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.

Result

When the task completes, the host is removed and the selected associated storage resources are released.

Related concepts

- About removing hosts and related resources, or vCenters on page 79

Registering file servers

This module describes how to register file servers.

About registering and removing file servers

You can register file servers to be managed by Device Manager so that the volumes of the storage system can be used by the file server.

Device Manager manages file servers based on information sent from NAS Platform family, NAS Platform F, or Data Ingestor.

Note: A system drive is the basic (logical) storage element managed by the NAS Platform family, and is equivalent to a storage system volume. Where volumes and system drives do not need to be differentiated, they are generically and collectively referred to as volumes.

When a file server is registered you can:

- Allocate and unallocate volumes
- Examine file server information

- Determine the usage status of storage resources from the file server point of view

From Device Manager, you can activate the file server management software, and change the settings of the file server or reference detailed file server information.

From Device Manager, in addition to the operations above, and depending on the versions of the file server management software that are linked and the file server type, file systems can be created, shares can be added, and file server reports and alerts can be viewed in the dashboard.

If the file server to be registered is NAS Platform v10.2.3071 or later, use Device Manager to set the Admin services EVS information (the IP address of Admin services EVS and the supervisor account information) after the registration. This enables file systems to be created on NAS Platform EVSs by using Device Manager.

If file servers become unnecessary for reasons such as reorganization, use file server management software to remove the settings for linking to Device Manager, then remove the file servers from the file server list by using HCS. The file servers are excluded from management by HCS, but configuration information, such as the file system and file sharing created on the file server, and paths set during volume allocation, are retained.

For details about the system requirements of the file server, see the *Hitachi Command Suite Administrator Guide*.

Related concepts

- About managing NAS Platform file systems and shares on page 246
- About managing NAS Platform F and Data Ingestor file systems and shares on page 273

Related tasks

- Registering file servers on page 84
- Changing the name of a file server on page 85
- Changing the Admin services EVS settings of NAS Platform on page 86
- Allocating volumes to selected file servers on page 206

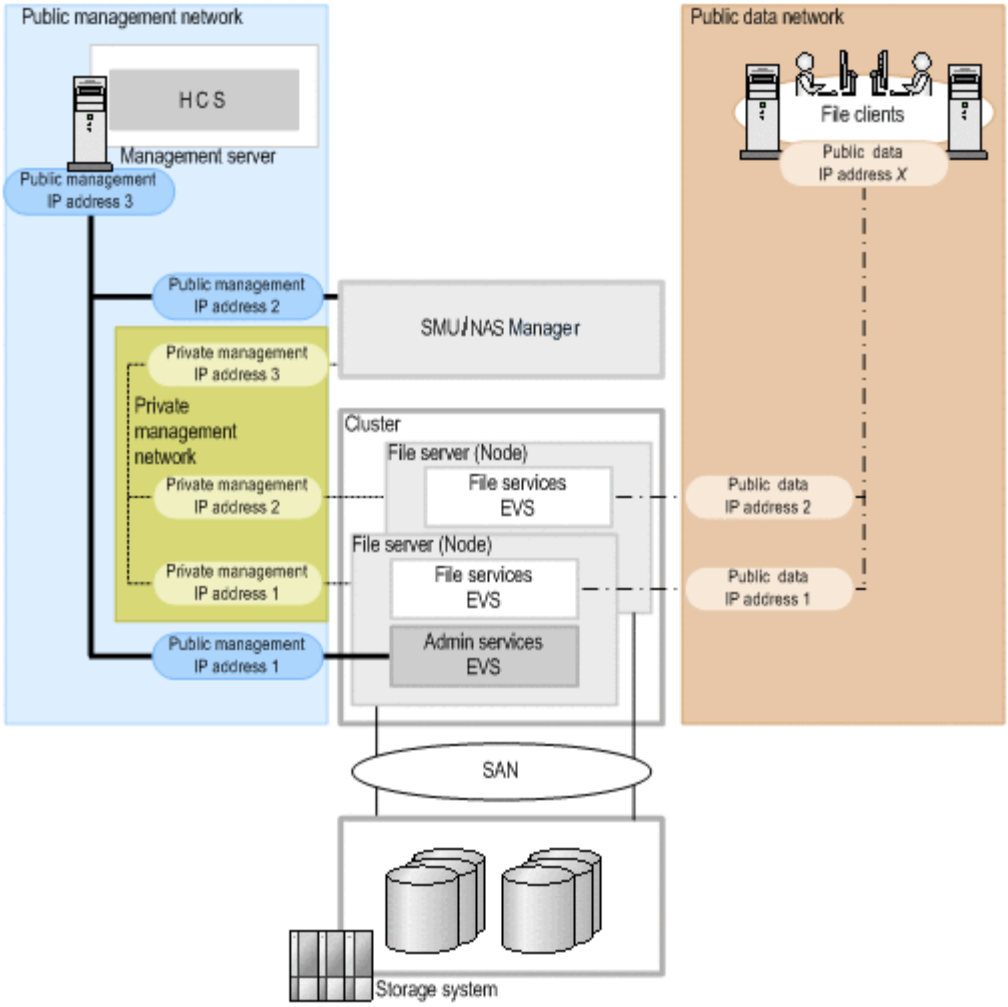
Related references

- Environment settings for NAS Platform family on page 82
- Workflow for providing NAS Platform file shares on page 250

Environment settings for NAS Platform family

To manage the NAS Platform using Device Manager, the management server must be in a system configuration in which communication with SMU/NAS Manager and Admin services EVS can be performed.

The following diagram illustrates the network relationships between the public management network and its connectivity with the NAS Platform over public IP addresses (SMU/NAS Manager and Admin services EVS). The storage system connectivity illustrates the use of storage system volumes by the file server, and File services EVS illustrates data services to clients over the public data network. See additional connectivity details below.



- Locate the management server on a network that can communicate with SMU/NAS Manager. To use the Device Manager GUI to create file systems for firmware version 10.2.3071 or later of NAS Platform, set up the system so that it can also communicate with Admin services EVS.
- Make sure that the version of SMU/NAS Manager is the same or later as the firmware version of the file server (node).
- Make sure that the versions of the firmware for the file servers (nodes) in a cluster are all the same.
- If Device Manager is managing firmware version 10.2.3071 or later of NAS Platform, register the following information for each cluster by using the Device Manager GUI:

- IP address of Admin services EVS (Public management IP address 1 in the figure).
  The IP address can be checked in the EVS Management page of SMU.
- User account for the Server Control (SSC).
  A supervisor account is set up as the default user.

Note: If you want to use the Device Manager GUI to check information such as the system drive and storage pool information of the file server, the file system information, or the sharing and exporting information, you need to configure LUN security for the file server from the storage system volume.

Related concepts

- About registering and removing file servers on page 81

Registering file servers

To register the file server (NAS Platform family, NAS Platform F, or Data Ingestor) as the resource managed by Device Manager, specify the Device Manager information by using the management software on the file server.

If you register NAS Platform version 10.2.3071 or later, specify the Admin services EVS information by using Device Manager. This establishes linkage between Device Manager and the Admin services EVS software on the NAS Platform.

Before you begin

- In Device Manager, register the storage system that the file server uses.
- Set up the file server management software:
  - System Management Unit (SMU) for NAS Platform family
  - Hitachi File Services Manager (HFSM) for NAS Platform F and Data Ingestor
- If the file server to be registered is NAS Platform v10.2.3071 or later, gather IP addresses for Admin services EVS (a virtual machine that manages the file server), and the user account for the Server Control (SSC).

Procedure

1. Register Device Manager information by using the file server management software.

   Device Manager users that are specified in the management software of the file server must belong to the **PeerGroup** as a user group. As a best practice, do not use the default **HaUser** user account. Instead, create a new account in Device Manager.
2. If the file server to be registered is NAS Platform v10.2.3071 or later, synchronize the information between the file server and Device Manager

using SMU/NAS Manager. Afterward, complete the following steps from Device Manager to set the Admin services EVS information:

a. On the **Administration** tab, select **Managed Resources**.
b. On the **File Servers** tab, select the NAS Platform to be registered, and click **Edit File Server**.
c. Enter the Admin services EVS information, and then submit the task.
d. From the **Data Collection Tasks** tab, verify the task status.

Result

File server information will be sent to Device Manager at regular intervals, and can be viewed in the File Servers tree on the Resources tab.

Tip: To remove registered file servers, use file server management software to remove the settings for linking to Device Manager. Then select the file servers to be removed from the HCS file server list in the Administration tab, and click Remove File Servers.

Related concepts

- About registering and removing file servers on page 81

Related references

- Priority for acquiring the WWN or iSCSI name on page 68
- Workflow for providing HNAS F and HDI file shares (HFSM) on page 275

Changing the name of a file server

You can change the file server name or IP address by changing the information registered in Device Manager, followed by changing the information on the file server.

Note: If the file server name or IP address on the file server side does not match the file server name (*file-server-name@IP-address*) registered in Device Manager, Device Manager cannot correctly recognize which file server information has been sent when the information is sent from a file server, and the same file server might be repeatedly registered under different names. Therefore, file server information might not be updated or an SNMP trap sent from the file server might not be displayed as an alert. If a file server is repeatedly registered, you need to remove the file server registered with the wrong name or an old name from Device Manager.

Procedure

1. On the **Administration** tab, select **Managed Resources**.
2. On the **File Servers** tab, select a file server whose name you want to change and click **Edit File Server**.
3. Enter the new name for the file server using the format *file-server-name@IP-address*.

For NAS Platform in a non-cluster configuration, if an IP address is not assigned to the node, enter the file server name in the following format:

- If an IPv4 address is assigned to Admin services EVS: *file-server-name*@*IP-address-of-Admin-services-EVS*.
- If an IPv4 address is not assigned to Admin services EVS: *file-server-name*.

4. On the file server side, change the file server settings on the target file server.
5. If you are using System Management Unit (SMU)/NAS Manager v10.2.3071 or later, then in SMU/NAS Manager, synchronize information between the file server and Device Manager.

Related concepts

- About registering and removing file servers on page 81

Changing the Admin services EVS settings of NAS Platform

If you change the Admin services EVS information (IP address, or user account for the Server Control (SSC)) of NAS Platform v10.2.3071 or later, you also need to change the Admin services EVS settings registered in Device Manager.

Procedure

1. Change the settings of Admin services EVS on the file server side.
2. On the **Administration** tab of Device Manager, select **Managed Resources**.
3. On the **File Servers** tab, select a file server and click **Edit File Server**.
4. Enter the changed Admin services EVS information, and then submit the task.
5. Verify the task status on the **Data Collection Tasks** tab.

Result

The file server status displayed in the list of file servers on the Administration tab is `Normal`.

Related concepts

- About registering and removing file servers on page 81

4

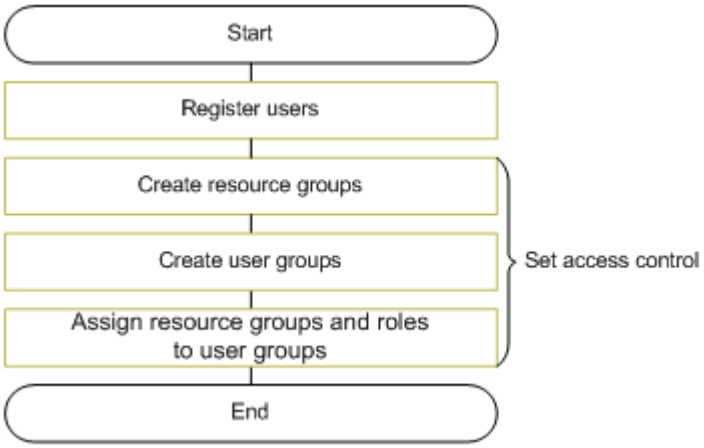
Setting up users and access control

This module describes how to control access to managed resources.

- ☐ Setting up users and access control
- ☐ About user accounts and controlling access to resources
- ☐ Creating and managing user accounts
- ☐ Controlling access to resources

Setting up users and access control

After users are registered, you can limit the scope of allowed operations for each user by configuring access control settings for users and storage resources.



To set access control you will need to create resource groups and user groups, then assign the resource groups and roles to the user groups.

Related concepts

- About user accounts and controlling access to resources on page 88
- About access control on page 99

About user accounts and controlling access to resources

Hitachi Command Suite provides built-in user accounts and the ability to add additional local user accounts and users from external authentication servers. You grant controlled access to storage resources by adding new users to user groups with assigned resource groups and roles (permissions). Built-in resource groups and user groups exist for administrative convenience.

The following two built-in user accounts are created when Hitachi Command Suite is installed. These accounts, and any additional user accounts you add, display in the list of user accounts.

- The System account (default password: `manager`) is a fully-privileged administrator account that is used to manage all HCS functionality, including HCS user accounts, user groups, and resource groups.
- The HaUser account (default password: `haset`) is the default user account used by Device Manager agents and exclusive account of management software for file servers or NAS Modules. The default role for the HaUser account is Peer and the PeerGroup is set for the HaUser account. The HaUser account belongs to PeerGroup as soon as the installation completes.

Log in using the System account to access user management functionality on the Administration tab to create local HCS user accounts. While creating user accounts, you can list available applications (as installed on the management server) and set user permissions for those applications. You must add users to at least one, or more, resource groups to determine the storage they can access. Together, application permissions and resource group/user group membership determine the scope of what each user can do in HCS.

When managing VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, or VSP F1500:

- For VSP G1000, VSP G1500, or VSP F1500, enable user authentication in HCS so that user accounts are authenticated when they log in to CCI and the SVP so that user accounts can be centrally managed.
- Secure communication settings:
  - For VSP Fx00 models or VSP Gx00 models, SSL/TLS communication must be enabled between the Device Manager server and SVP. Even if you do not configure the settings for SSL/TLS communication, when Device Manager is installed, the setting for encrypted communication is enabled by using the default certificate. Also, if a firewall is set up between the Device Manager server and the SVP, it might be necessary to register a firewall exception. For details about configuring separate SSL/TLS communication settings by using a different security policy, or setting up a firewall between the Device Manager server and the SVP, see the *Hitachi Command Suite Administrator Guide*.
  - For a VSP G1000, VSP G1500, or VSP F1500 storage system, SSL/TLS communication must be enabled between the Device Manager server and storage system for using HCS authentication for the SVP and CCI login. Even if you do not configure the settings for SSL/TLS communication, when Device Manager is installed, the setting for encrypted communication is enabled by using the default certificate. Also, if a firewall is set up between the Device Manager server and the SVP, it might be necessary to register a firewall exception. For details about configuring separate SSL/TLS communication settings by using a different security policy, or setting up a firewall between the Device Manager server and the storage system, see the *Hitachi Command Suite Administrator Guide*.
- User accounts should be created with user names and passwords compatible with HCS and VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, or VSP F1500 components.
- For VSP G1000, VSP G1500, or VSP F1500, if a user account that is used to perform operations by using CCI or the SVP is already registered in Hitachi Device Manager - Storage Navigator, also register that user account in HCS.
- Create an administrator user account for Hitachi Device Manager - Storage Navigator that can be used if HCS is not available.

Note: For VSP G1000, VSP G1500, or VSP F1500, if HCS authentication of user accounts is disabled when logging into CCI or SVP, you must specify the same user account information and access control to storage resources in both HCS and Hitachi Device Manager - Storage Navigator.

You can also manage user accounts by linking to an external authentication server, such as an LDAP directory server, RADIUS server, or Kerberos server. However, the built-in accounts (System and HaUser) cannot be authenticated on an external authentication server. The HCS user account used to connect to external authentication servers and external authorization servers is managed as a Windows Active Directory (authorization) group. Permissions that are specified for authorized groups are also applied to users who belong to nested groups.

Application permissions

After adding basic user information such as username, password, email, and description, set permissions for available applications, such as:

- Tier management (CLI)
- Replication management
- Performance management

Permissions include View, Execute, Modify, and Admin. These permissions control what the user can do on the related tabs, and possibly elsewhere.

By selecting the admin permission for the User Management application, users can assist in user management tasks, such as:

- Specifying user settings
- Creating user groups for Device Manager and Tiered Storage Manager
- Assigning resources and roles to user groups
- Reporting user, and user group information in CSV format
- Specifying security settings (such as locking an account)

Resource groups and User groups

All Resources is a built-in resource group (created by default) that contains the built-in user groups called AdminGroup, ModifyGroup, PeerGroup, and ViewGroup. Adding a user to ViewGroup allows that user to see all registered storage systems and related details such as parity groups. Adding the user to ModifyGroup enables task-related buttons and tasks listed under General Tasks, thus allowing the user to work with registered resources. Essentially, as a member of an All Resources group, you have access to the Device Manager GUI and Tiered Storage Manager GUI elements.

Each storage system has a resource group named Default ResourceGroup. If you have three registered storage systems you will see three of these groups listed in addition to All Resources. Default ResourceGroup is used to provide Admin, Modify, or View permissions (roles) to one or more users in a user-defined user group so they have access to the specific resources of the

storage system. In other words, instead of placing a user in the All Resources group, you can place them in one or more storage system resource groups to narrow the scope of what they can view or manage. To do this, you must create a named user-defined group and edit the resource group to add the user-defined group and one or more users, whose permissions (roles) can be set independently as you add them. Additionally, if the default resource group is for VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, or VSP F1500, you can select Custom roles which are more specific, such as roles for provisioning, or copy pair management and tasks. Multiple roles can be combined.

For very specific control over access to resources, consider creating user-defined resource groups. You can identify specific parity groups and LDEVs that members of your user-defined user group can access. As with default resource groups, for VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, or VSP F1500, you can select Custom roles.

Tip: For details about the required permissions for executing each command of the Tiered Storage Manager CLI, see the *Hitachi Command Suite Tiered Storage Manager CLI Reference Guide.*

Related tasks

- Creating a user account on page 91

Related references

- User ID and password policies on page 92

Creating and managing user accounts

Create user accounts and assign permissions.

Creating a user account

All users not allowed to log in with the System account require a user account for access to HCS.

A user account consists of general user profile information (User ID, Password, Full Name, E-mail, and Description).

Procedure

1. On the **Administration** tab, click **Users and Permissions**. This will launch a user management window.
2. Click **Users** to display the current user list.
3. Click **Add User** and specify user profile information.
4. Click **OK**.

Result

The user list is re-displayed and will include the new user.

Related tasks

- Editing the profile for a user account on page 93
- Deleting user accounts on page 98

Related references

- User ID and password policies on page 92

User ID and password policies

User IDs and passwords must adhere to specific requirements.

The User ID and password requirements for HCS, the SVP (VSP G1000, VSP G1500, VSP F1500), Device Manager - Storage Navigator, and Command Control Interface (CCI) vary.

When using HCS as an authentication server for VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, or VSP F1500 storage systems, User IDs and passwords must be valid for both HCS and the SVP (VSP G1000, VSP G1500, VSP F1500) or Device Manager - Storage Navigator, and for HCS and CCI. For VSP Fx00 models and VSP Gx00 models, you must use an HCS server as an authentication server.

Table 4 Login account requirements

Component	Item	Length	Requirements
HCS	User ID	1-256	A-Z, a-z, 0-9 ! # $ % & ' ( ) * + - . = @ \ ^ _ |
	Password	1-256	Same as above
SVP (VSP G1000, VSP G1500, VSP F1500)	User ID	1-128	Alphanumeric (ASCII code) characters "! # $ % & ' - . @ ^ _ `{ } ~
	Password	6-127	Alphanumeric (ASCII code) characters (All symbols)
Device Manager - Storage Navigator	User ID	1-256	Alphanumeric (ASCII code) characters # $ % & ' * + - . / = ? @ ^ _ `{|} ~
	Password	6-256	Alphanumeric (ASCII code) characters All symbols
CCI	User ID	1-63	Alphanumeric (ASCII code) characters - . / @ \ _
	Password	6-63	Alphanumeric (ASCII code) characters - , . / : @ \ _
NAS Manager	User ID	1-20	Alphanumeric (ASCII code) characters

Component	Item	Length	Requirements
			The following symbols: - . _ Note: Do not specify a user ID that only consists of periods (..) or that begins with a hyphen (-). If you specify such information, you will be unable to log in to NAS Manager.
	Password	6-256	Alphanumeric (ASCII code) characters All symbols

If using external authentication servers such as LDAP (and others), note the following:

- User IDs and passwords must be valid for the external authentication server and Hitachi Command Suite products.

A password policy can be configured from the Administration tab to enforce stronger passwords. If using external authentication, the password enforcement must be compatible.

Related concepts

- About user accounts and controlling access to resources on page 88

Related tasks

- Creating a user account on page 91
- Changing the password for a user account on page 94
- Changing your own password on page 94
- Configuring external authentication for users on page 96
- Configuring external authentication for groups on page 97

Editing the profile for a user account

Modify the name, email address, and description for a user account.

Procedure

1. On the **Administration** tab, click **Users and Permissions**.
   This will launch a user management window.
2. Click **Users**, select the target user by clicking the **User-ID** link, and click **Edit Profile**.
3. Edit the profile information for the user, and then click **OK**.
   The user profile is displayed.
4. Confirm the updated user profile information.

Related tasks

- Changing permissions for a user account on page 95

- Editing your own user profile on page 94

Editing your own user profile

As your user attributes change, you will need to update your user profile.

Procedure

1. On the **Administration** tab, click **User Profile**.
   Your user information is displayed.
2. Click **Edit Profile**.
3. Edit the profile information and click **OK**.
4. Confirm that the updated user profile information appears in the **Users** area.

Related tasks

- Changing your own password on page 94

Changing the password for a user account

As user passwords expire or are compromised, they can be changed.

Procedure

1. On the **Administration** tab, click **Users and Permissions**.
   This will launch a user management window.
2. Click **Users**, select the target user by clicking the **User-ID** link, and click **Change Password**.
3. Enter the new password and verify it.
4. Click **OK**.
5. Confirm that the user account can log in with the new password.

Related tasks

- Changing your own password on page 94

Changing your own password

As your password expires or is compromised, it will need to be changed.

Procedure

1. On the **Administration** tab, click **User Profile**.
   Your information is displayed.
2. Click **Change Password**.
3. Type the new password and verify it.
4. Click **OK**.
5. Log in with your new password.

Result

Your password is changed.

Related concepts

- About user accounts and controlling access to resources on page 88

Related tasks

- Changing the password for a user account on page 94

Related references

- User ID and password policies on page 92

Changing permissions for a user account

To grant a user new permissions or remove existing permissions, change permission settings in the user account.

Tip: For a user of Device Manager or Tiered Storage Manager (GUI), specify a role for the user group which is assigned to the user, instead of granting user permissions.

Procedure

1. On the **Administration** tab, click **Users and Permissions**.
   This will launch a user management window.
2. Click **Users**, select the target user by clicking the **User-ID** link, and click **Change Permission**.
3. Edit the permissions and click **OK**.
   The user account is re-displayed, including granted permission.
4. Verify the correct user permissions are selected.

Result

The user permissions are changed.

Related references

- Required roles and resource groups by function on page 114

Changing the lock status of user accounts

A user account can be locked or unlocked by an administrator.

Procedure

1. On the **Administration** tab, select **Users and Permissions**.
   This will launch a user management window.
2. Click **Users**, select the check box for the user whose lock status you want to change.
3. Click **Lock Users** or **Unlock Users**.

A verification dialog box displays.

4. Click **Ok** to lock or unlock the account, or click **Cancel**.
5. Verify that the user account has been locked (a lock icon displays in the user list), or that the previously locked user can now log in.

Related concepts

- About user accounts and controlling access to resources on page 88

Related tasks

- Changing the password for a user account on page 94

Configuring external authentication for users

External authentication systems can be used to authenticate user logins.

External authentication systems, such as LDAP (for example, Active Directory), RADIUS, or Kerberos may be used to authenticate HCS users as they log in. You can re-configure existing accounts, or create new accounts to use external authentication.

Before you begin

- The HCS server must be linked to an external authentication server. See the *Hitachi Command Suite Administrator Guide*.
- The HCS server must be configured to support user authentication, which activates the Change Auth button in the GUI, and which presents authentication options such as Internal for a local account, or LDAP for external authentication.
- The HCS user ID must exist on the external authentication server. It is recommended that user ID information be acquired from the external authentication server administrator before creating accounts.

Procedure

1. From the **Administration** tab, select **Users and Permissions**.
2. Select **Users** folder, then select one or more users (using the check box) whose authentication method you want to change, or click **Add User** to create a new account.

Note: When creating a new account, only the **User ID** is required for external authentication, and must match a user ID on the external authentication server. For a local (internal) account, a **User ID** and **Password** are both required. When external authentication is available, new user accounts created without a password value are automatically configured to use external authentication (for example, LDAP is selected for you). Fill in the desired fields, and click **OK** to create the user account.

If you are completing an external authentication in a multiple-domain configuration or by using realms, specify a user ID that includes the domain name or realm name for the **User ID**. Example: *user-name@domain-name* or *realm-name*.

When you log in, also specify your user ID in this format.

3. If you have selected existing users, click **Change Auth**. A dialog box is displayed. From the drop down list, select the desired authentication method (for example, LDAP) and click **OK**. The user list will be re-displayed.
4. Review the **Authentication** column to verify the authentication method.

Result

On the next login attempt by each user, the users login credentials (user ID and password) will be validated using the external authentication server.

Tip: Set permissions or roles so that the registered user can perform necessary operations using HCS products. Also consider adding user accounts to user groups with assigned roles for controlled access to resource groups.

Related concepts

- About user accounts and controlling access to resources on page 88

Related tasks

- Configuring external authentication for groups on page 97

Configuring external authentication for groups

External authentication systems can be used to authenticate user groups.

External authentication systems, such as LDAP (for example, Active Directory), RADIUS, or Kerberos may be used to authenticate HCS user group members as they log in. You can configure one or more user groups, from one or more external authentication servers.

When linking with an external authentication server, if using together with Active Directory as an external authorization server, user permissions can be managed by using the Active Directory groups (authorization groups) registered on the external authorization server. In this case, user permissions are specified for each group.

Before you begin

- The HCS server must be linked to an external authentication (authorization) server. See the *Hitachi Command Suite Administrator Guide*.
- The HCS server must be configured to support group authentication, which activates the Groups folder in the GUI.

- The HCS user group must exist on the external authentication (authorization) server. It is recommended that domain and group information, as required below, be acquired from the external authentication server administrator.

Procedure

1. From the **Administration** tab, select **Users and Permissions**.
2. Click the **Groups** folder to display the **Domain List**. This is a list of external authentication servers listed by domain name, and host name or IP address. If the **Groups** folder is not displayed, see the pre-requisites above.
3. Select the desired **Domain Name** to display the **Group List**, which may be empty ('No Groups' is displayed). Click **Add Groups**.
4. Enter the **Distinguished Name** for the group. Use **Check DN** to verify a correct DN entry. Click **Ok** to save your group and re-display the **Group List**. Note that the **Group Name** is derived from the entered DN. To specify multiple groups, note that:
   - You can add multiple DNs at the same time using the "**+**" button
   - If multiple DNs are listed, you can remove an entry with the "**-**" button
   - **Reset** clears all DN entries
5. From the **Group List**, click the **Group Name** link, then click **Change Permission** and set the HCS permissions for the group (repeat this for each new group).
6. Your groups will now be visible from the **Administration** tab, **User Groups**. You can affiliate the groups with resource groups and roles, just like HCS user groups. If you delete external authentication groups from **Users and Permissions** at a later time, the groups are also removed from the **User Groups** list.

Result

On the next login attempt by each group member, the users login credentials (User ID and Password) will be validated using the external authentication (authorization) server.



Tip: To delete registered authorization groups, select the check boxes of the groups to be deleted, and then click **Delete Groups**.

Related concepts

- About user accounts and controlling access to resources on page 88

Related tasks

- Configuring external authentication for users on page 96

Deleting user accounts

If user accounts are no longer needed for accessing HCS, for example if users leave the organization, you can delete the user accounts.

Procedure

1. On the **Administration** tab, select **Users and Permissions**.
2. Select **Users** in the navigation pane, and then select the users to delete.
3. Click **Delete Users**.
4. Click **OK**.

Result

The user accounts you deleted no longer appear in the list of user accounts.

Controlling access to resources

This module describes how to control access to resources.

About access control

Create an access control policy to allow secure data handling in multi-tenant environments and support more efficient and secure operations.

An access control policy is used for:

- Data center hosting services
- Management of departments in an organization
- Management of locations in an organization

To create an access control policy in a managed SAN environment, you create user accounts, add them to user groups, and then affiliate the user groups with resource groups and assigned roles to provide controlled access to functionality available in Device Manager and Tiered Storage Manager (GUI).

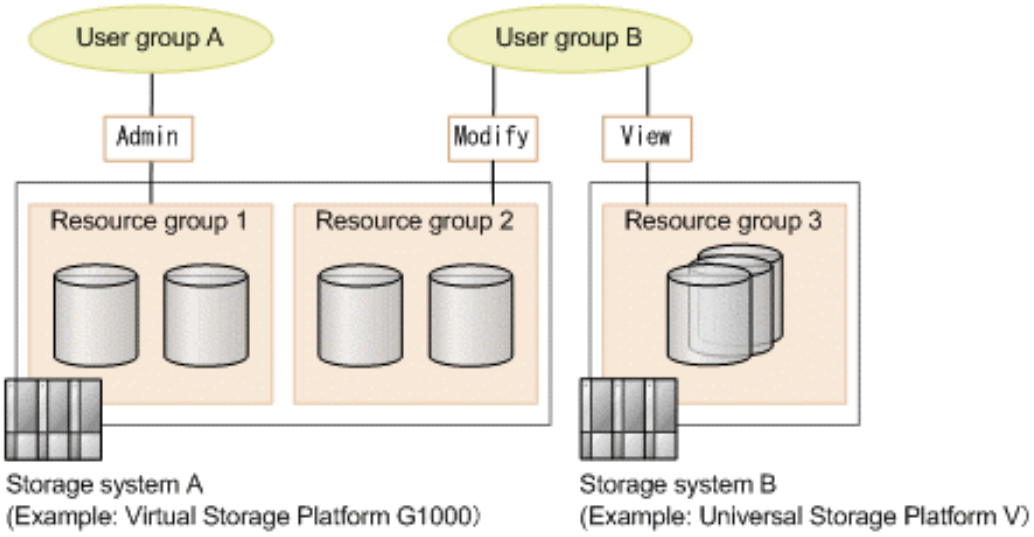
A user group, which is made up of local user accounts or accounts from external authentication systems, is a group of users who can access the same resources with the same user permissions. A resource group consists of storage system resources (storage systems, parity groups, DP pools, LDEV IDs, and storage ports). Assigned roles for resource groups provide either full, partial, or read-only access to resource group resources. When you assign resource groups and roles (user permissions, such as Admin, Modify, View or Custom) to a user group, resources are consistently controlled for the users in that group.

For VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, or VSP F1500 storage systems, you can use custom roles to specify one or more roles and user permissions at a more detailed, granular level. For example, you can allow:

- Provisioning operations
- Remote copy operations
- System resource operations
- Storage encryption key and authentication management
- Audit log management

Resource groups can be created in this configuration only for VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, VSP F1500, Virtual Storage Platform, or Unified Storage VM.

The following figure illustrates user groups and their permissions (standard Admin, Modify and View roles) for accessing resources. The use of custom roles is not shown here, but is illustrated in the user group topics. Custom roles provide more granular permissions to specific functionality.



For VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, VSP F1500, Virtual Storage Platform, or Unified Storage VM systems, physical configurations such as parity groups, and logical configurations such as LDEV IDs, are used to create resource groups. After resource groups are created, they can then be assigned to user groups.

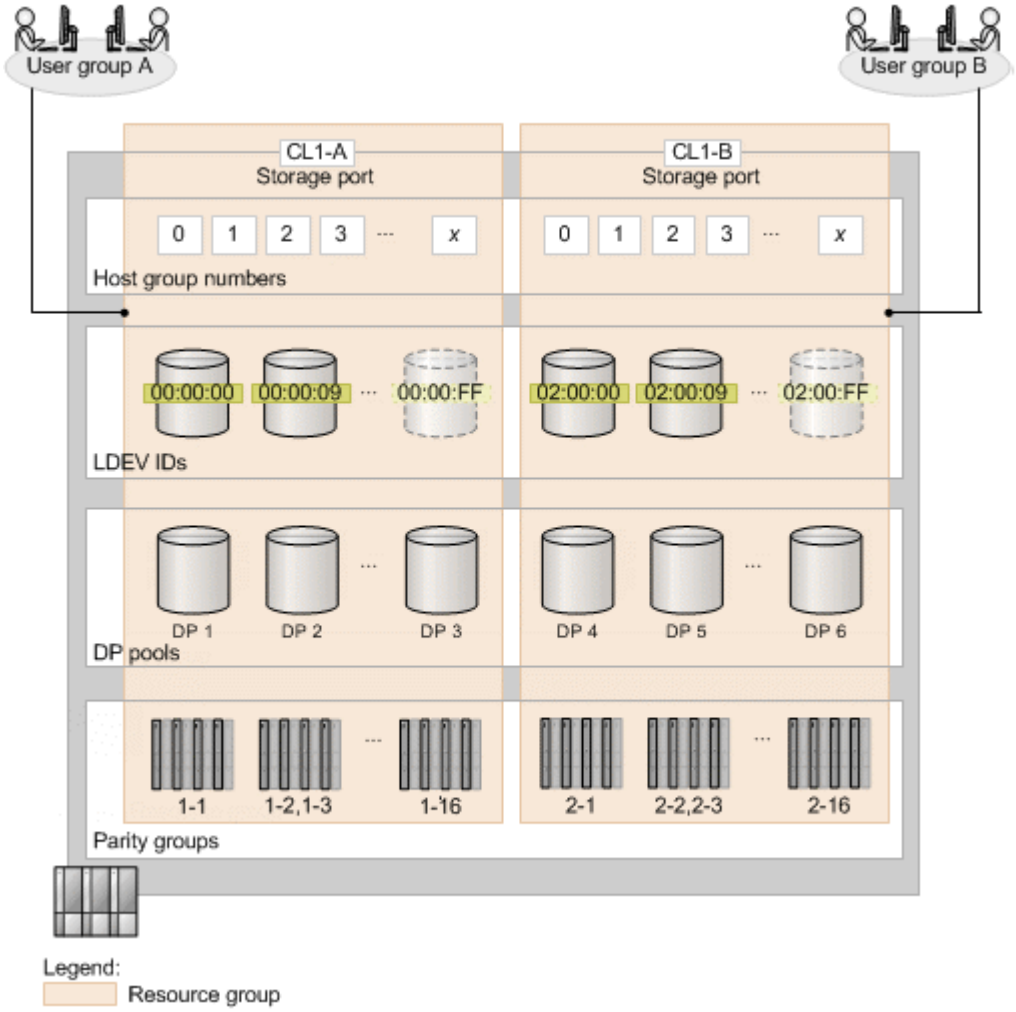
Related references

- Access control examples on page 100

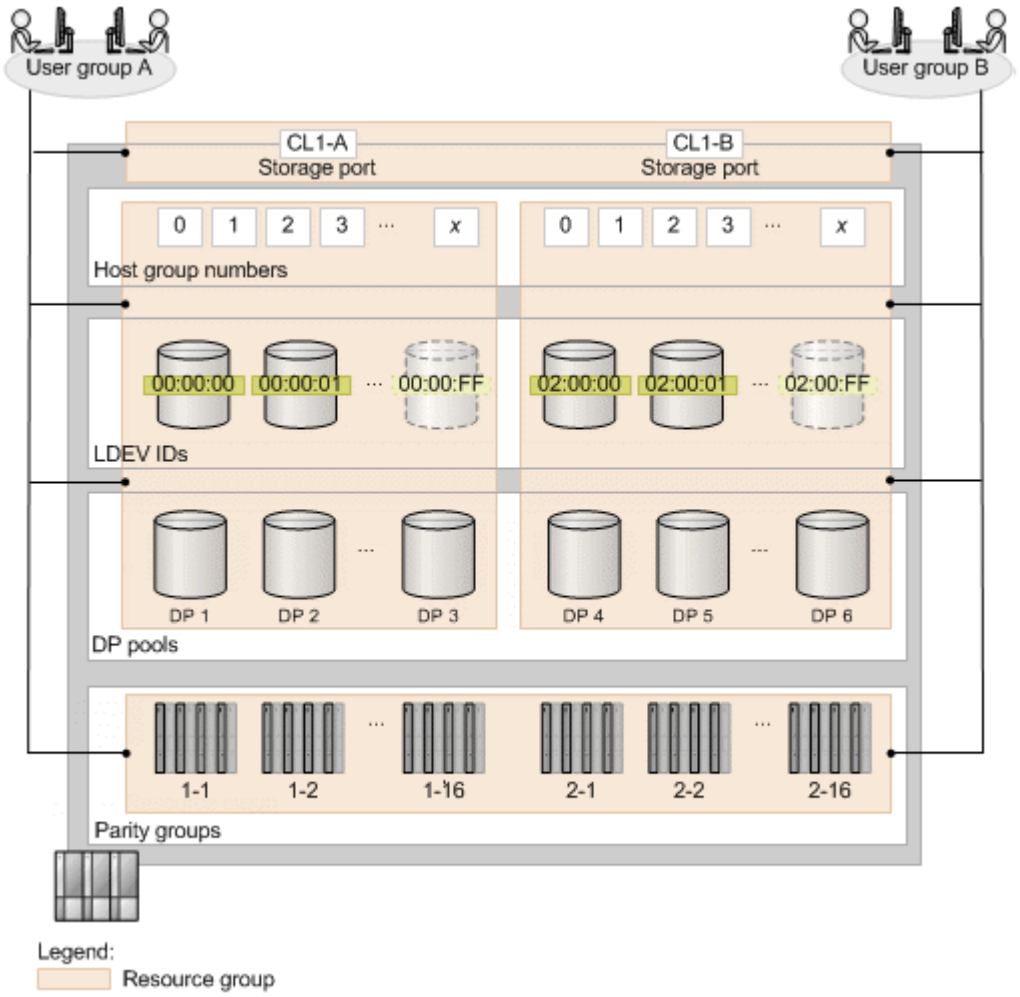
Access control examples

The following examples show how resource groups can control access in VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, VSP F1500, Virtual Storage Platform, or Unified Storage VM storage systems.

One method for dividing resources is to separate resources based on company location. For example, if you create resource groups based on location, the administrators in each location are limited to using only the resources that have been assigned to them, and are restricted from accessing the resources of other locations.



It is also possible to share physical resources (such as parity groups or storage ports) among departments, and divide only logical resources (such as DP pools, LDEV IDs, or host group numbers) by department. For example, you can assign resource groups that contain shared physical resources to all departments, and then assign individual resource groups that contain specific logical resources to the appropriate departments. This allows department administrators to use only the resources assigned to them, while still allowing for effective sharing of physical resources.



Related concepts

- About resource groups on page 102
- About user groups on page 107
- About access control on page 99

About resource groups

Resources can be grouped by system resource types that include storage system, parity groups, DP pools, LDEV IDs, storage ports, and host group/ iSCSI target numbers.

Note: When DP pools are registered to resource groups, related DP pool volumes and their LDEV IDs are also registered.

There are several types of resource groups:

- **All Resources** is a resource group that is created during management server installation and includes all resources managed by HCS. For example, a user who is a member of one of the built-in user groups for All Resources has access to all storage systems.

- **Default ResourceGroup** is the name for default resource groups that are created for a storage system when it is discovered and registered. A user who is a member of a user group in a default resource group has access to all of the storage system resources.
- **Resource pool** is another type of resource group. A resource pool is a resource group to which resources of a virtual storage machine belong in VSP G1000, VSP G1500, VSP F1500, VSP Fx00 models and VSP Gx00 models, when the resources have not been added to any individual resource group. There are two types of resource pools. There are resource pools on the default virtual storage machine, and resource pools that are automatically created on user-defined virtual storage machines. You can check the resource pools on user-defined virtual storage machines from the resource group list.
- Resource groups for System LUs in VSP G400, G600, G800, and VSP Fx00 models with NAS modules are automatically created.
- User-defined resource groups defining more specific storage access can be created for VSP Fx00 models, VSP Gx00 models (excluding NAS modules), VSP G1000, VSP G1500, VSP F1500, Virtual Storage Platform, and Unified Storage VM depending on the operating environment. Resources can be grouped by parity groups, DP pools, LDEV IDs, storage ports, or host group/iSCSI target numbers. Resource group definitions in Device Manager are applied to VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, or VSP F1500 storage systems. However, these resource group definitions are not applied to other storage systems.

Each resource is automatically registered in the All Resources and Default resource groups created for its storage system (this group cannot be deleted). If a volume that is part of a LUSE volume is registered in a resource group, other volumes in that LUSE volume are also registered in the same resource group.

For VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, or VSP F1500, when you register a part of a parity group that is part of a concatenated parity group to a resource group, other parity groups that are a part of the concatenated parity group will also be registered in the same resource group automatically. If the resource is in a VSP Fx00 models, VSP Gx00 models, Virtual Storage Platform, VSP G1000, VSP G1500, VSP F1500 or Unified Storage VM system, you can register it in only one user-defined resource group.

Related concepts

- About user groups on page 107
- About virtual storage machines on page 322

Related tasks

- Creating resource groups on page 104
- Editing a resource group on page 105

- Assigning resource groups and roles to a user group on page 136

Related references

- Prerequisites for creating resource groups on page 104
- Access control examples on page 100

Prerequisites for creating resource groups

Resources can be grouped by system resource types that include storage systems, parity groups, DP pools, LDEV IDs, storage ports, and host group/iSCSI target numbers.

The following list identifies the conditions for creating a user-defined user group for VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, VSP F1500 Virtual Storage Platform or Unified Storage VM.

- **Parity Groups:** Includes parity groups and volumes in external storage systems. Users with Modify, Storage Administrator (Provisioning), or higher roles for parity groups, or the LDEV ID of a DP pool volume, and an unused LDEV ID is assigned to the user, the user can create a volume. For VSP Fx00 models, VSP Gx00 models and VSP G1000, VSP G1500, and VSP F1500 storage systems, when you register a part of a parity group that is part of a concatenated parity group to a resource group, other parity groups that are a part of the concatenated parity group will also be registered in the same resource group automatically.
- **DP Pools:** Includes DP pools consisting of DP pool volumes with LDEV IDs.
- **LDEV IDs:** Includes parity groups and volumes in external storage systems. Non-existent IDs can also be specified. Users with Modify, Storage Administrator (Provisioning), or higher roles for parity groups or DP pools and assigned an unused volume ID, can create a volume.
- **Storage Ports:** Users with Modify, Storage Administrator (Provisioning), or higher roles for ports and assigned an unused Host Group Number can create a host group that has that host group number.
- **Host Group/iSCSI Target Numbers:** Non-existent numbers can also be specified. Users with Modify, Storage Administrator (Provisioning), or higher roles for ports and assigned an unused Host Group Number can create a host group or iSCSI target.

Related concepts

- About resource groups on page 102

Related tasks

- Creating resource groups on page 104

Creating resource groups

User created resource groups can be used to group system resource types including storage systems, parity groups, DP pools, LDEV IDs, storage ports, and host group/iSCSI target numbers.

Resource groups, which are user defined, can be created for the following storage systems:

- VSP Gx00 models
- VSP Fx00 models
- VSP G1000
- VSP G1500
- VSP F1500
- Virtual Storage Platform
- Unified Storage VM

Procedure

1. On the **Administration** tab, in the **Administration** pane, select **Resource Groups**.
2. Click **Create Resource Group**.
3. Enter a name and description, and select the storage system providing the resources.
4. Using the tabs, specify the parity groups, DP pools, LDEVs, ports, host groups, iSCSI targets (or a mix of resources) for the resource group.
5. Click **Submit** to register this as a task.
6. You can check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.

Result

The new resource group is displayed, and can be assigned to an existing user group using the Edit User Group button. You can also assign resource groups when creating new user groups with Create User Group.

Related tasks

- Editing a resource group on page 105
- Assigning resource groups and roles to a user group on page 136
- Deleting resource groups on page 106

Related references

- Access control examples on page 100
- Prerequisites for creating resource groups on page 104

Editing a resource group

You can edit storage system resources in an existing resource group.

Information about resource groups can be modified to reflect changing access control requirements.

Procedure

1. On the **Administration** tab, in the **Administration** pane, select **Resource Groups**.
2. To edit a resource group, do one of the following:
   - From the **Resource Groups** pane, select the resource group, and click **Edit Resource Group**.
   - Click the resource group link, click **Actions** and select **Edit Resource Group**.
3. You can modify the resource group name and description, but not the storage system.
4. Modify the parity groups, DP pools, LDEVs, ports, iSCSI targets, or host groups to reflect your access control requirements.

   Note: To add or delete DP pool volumes, you must add or delete DP pools.

5. Click **Submit** to register this as a task.
6. You can check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.

Result

Depending on how you initiated your edit (see step 2), the resource group is displayed and you can confirm your changes, or you will be in the Resource Groups pane and can click the resource group link to confirm your changes.

Related tasks

- Creating resource groups on page 104
- Assigning resource groups and roles to a user group on page 136

Related references

- Access control examples on page 100

Deleting resource groups

If resource groups are no longer needed, you can delete the resource groups.

Procedure

1. On the **Administration** tab, select **Resource Groups**.
2. Select the resource groups to delete.

   The storage system default resource groups, All Resources, and resource pools cannot be deleted.
3. Click **Delete Resource Groups**.
4. Click **Submit**.

Result

The resource groups you deleted no longer appear in the list of resource groups.

About user groups

A user group consists of one or more users having the same permissions (role) for the same resources. An external authentication group can also be used as a user group. There are also built-in resource and user groups for administrative convenience.

For a user group, one or more resource groups are added, and a role assigned for each resource group. The types of roles are:

- Admin
- Modify
- View
- Custom

User group members are able to work with each resource group according to the assigned role (permissions) for the resource group. For example, a user group member with View access to a resource group can monitor, but not change the resource. Users can belong to multiple user groups, each with assigned resource groups and roles. Resource groups can be registered to multiple user groups.

For hosts and volumes that are managed as logical groups that correspond to businesses or organizations, when the logical groups are registered as private logical groups, only users who belong to the same user group will be able to use the logical groups.

The following default (built-in) user groups are assigned to the **All Resources** resource group (also built-in):

- AdminGroup (role: Admin and the permission for creating resource groups)
- ModifyGroup (role: Modify)
- ViewGroup (role: View)
- PeerGroup (role: Peer. This user group cannot be assigned to a resource group)

Note: If Compute Systems Manager v8.1 or later is installed on the HCS management server with Device Manager, the following user groups are also created:

- HCSM_AdminGroup
- HCSM_ModifyGroup
- HCSM_ViewGroup

If Hitachi Automation Director v8.1.1 or later is installed on the HCS management server with Device Manager, the following user groups unique to Automation Director are created:

- SubmitGroup
- DevelopGroup

Two special case user group assignments exist:

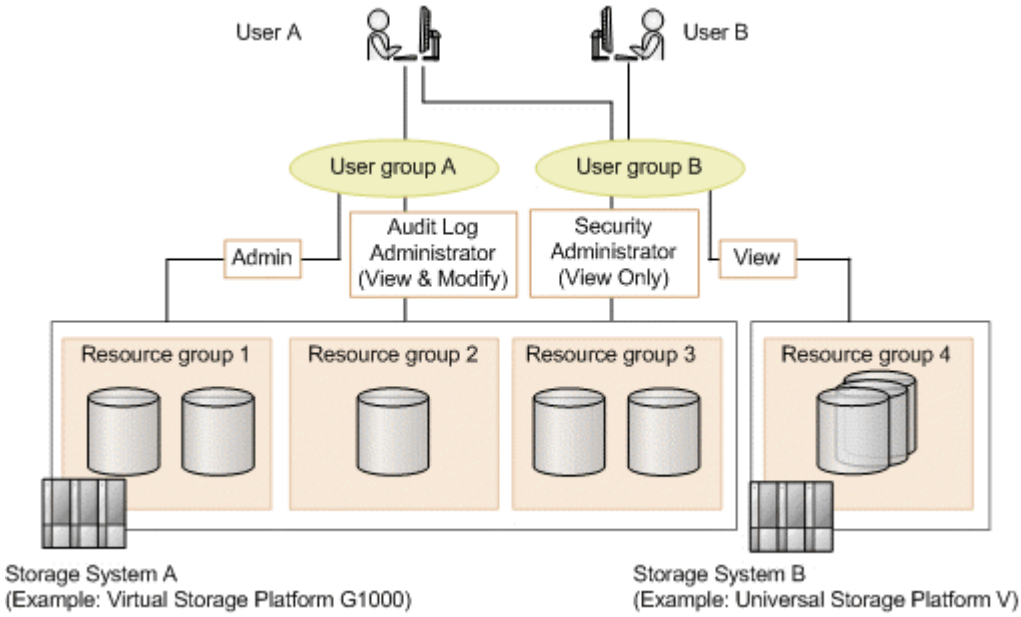
- The built-in account (user ID: HaUser) used by Device Manager agents and file servers or NAS modules is set to the PeerGroup immediately after the installation is completed, but can be set to another group later. To assign the Peer role to a user, register the user in PeerGroup.
- Authorized groups that have been registered to Hitachi Command Suite products can be used as user groups. Roles assigned to authorized groups are also applied to users who belong to nested groups.

For VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, VSP F1500, Virtual Storage Platform, or Unified Storage VM storage systems, if different roles are set as follows, the role set for each resource group is applied to all resource groups within the same storage system.

- When multiple resource groups in the same storage system are assigned to one user group, and a different role has been set for each resource group.
- When a user belongs to multiple user groups, and a different role has been set for the resource groups in the same storage system.

If the storage system is not VSP G1000, VSP G1500, VSP F1500, Virtual Storage Platform, Unified Storage VM, VSP Fx00 models or VSP Gx00 models, then the previous scenario does not apply. For example, in the following figure, User A and User B can access each resource group (RG) with the following roles, respectively.

User A can access RG1, RG2, and RG3 with the Admin, Audit Log Administrator (View & Modify) and Security Administrator (View Only) roles. User B can access RG3 with the Security Administrator (View & Modify) role, and access RG4 with the View role.



Some special cases apply:

- If a user has the Storage Administrator (Provisioning), Modify, or higher roles for parity groups or the LDEV ID of a DP pool volume, and an unused LDEV ID is assigned to this user, they can create a volume.
- If a user has the Storage Administrator (Provisioning), Modify, or higher roles for ports, and an unused Host Group ID is assigned to this user, they can allocate new volumes by using that Host Group ID.
- If the LDEV ID of a DP volume is assigned to a user, this user can view the DP pool to which the DP volume belongs and the DP pool volumes that compose the DP pool. If the LDEV ID of a DP pool volume is assigned to this user, they can view the pool to which the DP pool volume belongs.
- If a parity group is assigned to a user, this user can view all volumes that belong to the parity group from a list of volumes that appears when displaying the parity group information. If a parity group is not assigned to a user and only the LDEV IDs of the volumes belonging to the parity group are assigned to this user, they cannot view that parity group.

Note: The roles above determine the operation permissions of Device Manager and the Tiered Storage Manager GUI. To grant permissions to users in a user group for operating Tiered Storage Manager CLI, assign the user group to All Resources resource group and assign Device Manager roles, and then set the Tiered Storage Manager permissions required to execute each command. For details about the Device Manager roles and the permissions required to execute each command, see the *Hitachi Command Suite Tiered Storage Manager CLI Reference Guide*.

Related concepts

- About access control on page 99

User group roles

In Device Manager and Tiered Storage Manager (GUI), permissions are granted by assigning resource groups and roles to users in a user group.

For other HCS products, permissions are granted by setting permissions for each user, which, for example, allows you to grant permissions for the Device Manager GUI and CLI operations and for the Tiered Storage Manager GUI. To grant permissions to users in a user group for operating Tiered Storage Manager CLI, assign the user group to All Resources resource group and assign Device Manager roles, and then set the Tiered Storage Manager permissions required to execute each command. For details about the Device Manager roles and the permissions required to execute each command, see the *Hitachi Command Suite Tiered Storage Manager CLI Reference Guide*.

By specifying roles, resources are displayed for resources that belong to a resource group for which a user has permissions. The user can perform operations or reference information for the displayed resources.

Roles can be set for an external authentication group, just like for other user groups, when the external authentication group is used as a user group. By default, the View role is set for All Resources.

The following table describes roles and the tasks that can be performed when those roles are assigned.

Table 5 User permissions by role

Role	Device Manager Tasks	Tiered Storage Manager Tasks
Admin	The user can register resources to be managed, change settings, and view information. If the user is assigned to All Resources, the user can manage resource groups.	The user can create, edit, and delete tiers, perform operations from the Mobility tab, and perform migration tasks.
Modify	The user can register resources to be managed, change settings, and view information.	The user can create, edit, and delete tiers, perform operations from the Mobility tab, and perform migration tasks.
View	The user can view (reference) managed resources.	The user can view (reference) information about tiers, information in the Mobility tab, and list migration tasks.
Peer	This role applies only to Device Manager agents and file servers and cannot be assigned to resource groups and cannot be used to log in to HCS products. The Peer role cannot be assigned in combination with any permissions other than the User Management permissions.	Not applicable.
Custom	For VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, and VSP F1500, more granular roles, called custom roles, are available. The Admin, Modify, and View roles are	

Role	Device Manager Tasks	Tiered Storage Manager Tasks
	broad in scope, while custom roles are more specific. When selecting permissions for a user group associated with a default user group or user-defined resource group, multiple custom roles can be selected in combination to determine user capabilities. For users assigned to an All Resources built-in group, custom roles are not available as the built-in groups grant Admin, Modify, or View permissions only.	

Related references

- Custom roles on page 111
- Required roles and resource groups by function on page 114

Custom roles

Custom roles provide granular permissions for performing general HCS tasks, as well as additional tasks specific to VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, and VSP F1500. Available custom roles include Storage, Security, Audit Log, and Support roles.

The table below describes additional tasks (functions) and the required custom roles when selecting System GUI from menus or application panes to open Device Manager - Storage Navigator.

Custom roles must be assigned to resource groups with users to use them. The following custom roles can be assigned to both the default resource group for broad access to storage resources, and to user-defined resource groups for specific access to storage resources:

- Storage Administrator (Provisioning)
- Storage Administrator (Performance Management)
- Storage Administrator (Local Copy)
- Storage Administrator (Remote Copy)

Storage, security, and audit log custom roles, as well as the Support Personnel and User Maintenance roles for the VSP Fx00 models and VSP Gx00 models, that are not in the list above are generally for tasks concerning the storage system as a whole, such as security and auditing. These roles are assigned to the default resource group only.

Note: Custom roles cannot be assigned to users in the All Resources built-in resource groups as these groups permit View, Modify, or Admin permissions only.

Table 6 Custom roles

Custom role (permission)	Functions
Storage Administrator (Provisioning)[1]	Allows provisioning related operations: • Configuring caches • Configuring LDEVs, pools, and virtual volumes • Formatting and shredding LDEVs

Custom role (permission)	Functions
	• Configuring external volumes • Creating and deleting quorum disks used in a global-active device environment • Configuring alias volumes for Compatible PAV (VSP G1000, VSP G1500, and VSP F1500 only) • Creating and deleting global-active device pairs • Configuring Dynamic Provisioning • Configuring host groups, paths, and WWN/iSCSI names • Configuring Volume Migration except splitting Volume Migration pairs when using CCI • Configuring access attributes for LDEVs • Configuring LUN security
Storage Administrator (Performance Management)[1]	Allows performance monitoring: • Configuring monitoring • Starting and stopping monitoring
Storage Administrator (Local Copy)[1]	Allows pair operations for local copy: • Performing pair operations for local copy • Configuring environmental settings for local copy • Splitting Volume Migration pairs when using CCI
Storage Administrator (Remote Copy)[1]	Allows remote copy operations: • Remote copy operations in general • Managing global-active device pairs, except for creation and deletion
Storage Administrator (Initial Configuration)[1]	Allows initial configuration of storage systems: • Configuring settings for storage systems • Configuring settings for SNMP • Configuring settings for e-mail notification • Configuring settings for license keys • Viewing, deleting, and downloading storage configuration reports • Acquiring all the information about the storage system and refreshing
Storage Administrator (System Resource Management)[1]	Allows configuring various storage system resources: • Configuring settings for CLPR • Configuring settings for MP Blade or MP Unit • Deleting tasks and releasing exclusive locks of resources • Completing SIMs (VSP G1000, VSP G1500, and VSP F1500 only) • Configuring attributes for ports (VSP G1000, VSP G1500, and VSP F1500 only) • Configuring LUN security • Configuring Server Priority Manager • Configuring tiering policies
Security Administrator (View & Modify)	For global-active device: • Setting the reserved attribute for a volume to be used in a global-active device pair With the exception of user management, allows management of encryption keys and authentication for storage systems: • Creating an encryption key, configuring encryption

Custom role (permission)	Functions
	• Viewing and switching the location at which to create an encryption key • Backing up and restoring an encryption key • Deleting an encryption key that is backed up on the key management server, or key service processor • Viewing and changing the password policies for backing up an encryption key on the management client • Configuring the certificate used for SSL communication on the management client[2]. • Configuring the Fibre Channel authentication (FC-SP)
Security Administrator (View Only)	Allows viewing of storage system encryption keys and authentication settings: • Viewing information about encryption settings • Viewing information about encryption keys on the key management server, or key service processor
Audit Log Administrator (View & Modify)	Allows management of storage system audit logs: • Configuring audit log settings • Downloading audit logs
Audit Log Administrator (View Only)	Allows viewing of audit log settings for storage systems and downloading of audit logs: • Viewing storage audit log settings • Downloading audit logs
Support Personnel[3],[4]	Allows maintenance of the storage system. This role is for the exclusive use of service representatives.
User Maintenance	This role is for the exclusive use of service representatives. Allows storage system management, such as: • Viewing the storage system status • Installing OS security patches • Updating OSes • Performing basic maintenance tasks

Notes:

1. Custom roles also apply to general tasks performed on VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, and the VSP F1500 storage systems, such as:
- Refreshing storage system information
- Registering storage systems and hosts
- Managing tasks, logical groups, and storage tiers
- Displaying information
- Downloading components

2. For VSP G1000, VSP G1500, and VSP F1500, when a user account for logging in to the SVP or Command Control Interface (CCI) is authenticated by HCS, if the user account created in HCS is assigned the Security Administrator (View & Modify) role, that user account can be used to open the Tool Panel and configure the certificate. For details about this procedure, see the *System Administrator Guide* for your storage system.

3. For VSP G1000, VSP G1500, and VSP F1500, when a user account for logging in to the SVP or Command Control Interface (CCI) is authenticated by HCS, if the user account created in HCS is assigned the Support Personnel role, that user account can be used to log in to the SVP and perform tasks.

4. For VSP G1000, VSP G1500, and VSP F1500, when a user account for logging in to the SVP or Command Control Interface (CCI) is authenticated by HCS, if the user account created in HCS is assigned the Support Personnel role, that user account can be used to open the Tool Panel

Custom role (permission)	Functions
	and download dump files. For details about this procedure, see the *System Administrator Guide* for your storage system.

Related concepts

- About access control on page 99

Related references

- User group roles on page 110
- Required roles and resource groups by function on page 114

Required roles and resource groups by function

The following tables show the resource groups and roles that are required to perform each function in Device Manager or Tiered Storage Manager.

The first table below lists HCS functions, and the required resource groups and roles to perform the function.

The second table lists additional HCS functions for VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, and VSP F1500 and the required custom roles or roles to perform the functions.

Note: This topic describes only tasks that can be performed by using the GUI. For tasks that can be performed by using the CLI, see the *Hitachi Command Suite CLI Reference Guide* and *Hitachi Command Suite Tiered Storage Manager CLI Reference Guide*.

The following table groups related or similar functions.

Note: For custom roles, a hyphen (-) indicates that the task (function) cannot be performed with a custom role.

Table 7 Required resource groups and roles for performing functions

Function	Resource group	Required roles	
		Admin, Modify, View	Custom (VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, and VSP F1500)
Access control (Administration tab, resource groups) (Resources tab, logical groups)			
Assign resources and roles to user groups	All Resources	Admin You must have User Management Admin permission.	-
Create, delete, or edit resource groups	All Resources	Admin	-

Function	Resource group	Required roles	
		Admin, Modify, View	Custom (VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, and VSP F1500)
Create, edit, or delete public logical groups	Any	Admin or Modify	One of the following: Storage Administrator (Provisioning) Storage Administrator (Performance Management) Storage Administrator (Local Copy) Storage Administrator (Remote Copy) Storage Administrator (Initial Configuration) Storage Administrator (System Resource Management)
Create, edit, or delete private logical groups	Any	Any	
Downloads (Tools menu)			
Download related programs	Any	Admin or Modify	One of the following: Storage Administrator (Provisioning) Storage Administrator (Performance Management) Storage Administrator (Local Copy) Storage Administrator (Remote Copy) Storage Administrator (Initial Configuration) Storage Administrator (System Resource Management)
Link and launch			
Launch other HCS products	Any When starting Element Manager, the resource group to which the target resource belongs	Any	
Storage systems (Resources and Administration tabs)			

Function	Resource group	Required roles	
		Admin, Modify, View	Custom (VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, and VSP F1500)
Add storage systems Delete storage systems	All Resources	Admin	-
Edit storage system information (storage system name, IP address, host name, user name, or password)	All Resources	Admin or Modify	-
Apply storage system resource labels to Device Manager	Resource group to which the resources of the target system belong	Admin or Modify	-
Refresh storage systems	Resource group to which the resources of the target system belong	Admin or Modify	One of the following: Storage Administrator (Provisioning) Storage Administrator (Performance Management) Storage Administrator (Local Copy) Storage Administrator (Remote Copy) Storage Administrator (Initial Configuration) Storage Administrator (System Resource Management)
Acquire performance information when the storage systems are updated	Resource group to which the resources of the target system belong	Admin or Modify	Storage Administrator (Performance Management)
Hosts and vCenters (Resources and Administration tabs)			
Add hosts Edit hosts Refresh hosts Refresh vCenters Delete host (when deleting a host only) Resynchronize with hosts managed by Compute Systems Manager.	Any	Admin or Modify	One of the following: Storage Administrator (Provisioning) Storage Administrator (Performance Management) Storage Administrator (Local Copy) Storage Administrator (Remote Copy)

Function	Resource group	Required roles	
		Admin, Modify, View	Custom (VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, and VSP F1500)
			Storage Administrator (Initial Configuration) Storage Administrator (System Resource Management)
Delete hosts (when deleting a host and its related storage resources)	Resource group to which the storage resource to be deleted belongs	Admin or Modify	Storage Administrator (Provisioning)
Scan hosts Merge hosts	All Resources	Admin or Modify	-
LUN paths, HBAs, host modes (volume allocation, edit LUN paths, host HBA tasks)			
Edit WWN nickname Edit LUN path, edit host mode or host mode option, or edit LUN path when the HBA is replaced	Resource group to which the target resource belongs	Admin or Modify	Storage Administrator (Provisioning)
Data collection tasks (Administration tab)			
View Data Collection tasks list	Any	Any	
Delete or view details of data collection tasks (add or delete storage systems)	All Resources You can perform actions only on tasks that you have created even if All Resources is not allocated.	Admin	-
Delete or view details of data collection tasks (edit storage systems)	All Resources You can perform actions on tasks that you have created even if All Resources is not allocated.	Admin or Modify	-
View details of data collection tasks (refresh storage systems or update database)	All Resources or resource groups to which a target storage system belongs You can perform actions on tasks that you have created even if All Resources is not allocated. For VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, and VSP F1500, the default	Admin or Modify	Storage Administrator (System Resource Management)

Function	Resource group	Required roles	
		Admin, Modify, View	Custom (VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, and VSP F1500)
	resource group can also perform this action.		
Delete data collection tasks (refresh storage systems or update database)	All Resources You can perform actions on tasks that you have created even if All Resources is not allocated. For VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, and VSP F1500, the default resource group can also perform this action.	Admin or Modify	Storage Administrator (System Resource Management)
View details of data collection tasks (add, edit, update a host, or delete only hosts) View details of data collection tasks (tasks related to a file server)	Any	Admin or Modify	One of the following: Storage Administrator (Provisioning) Storage Administrator (Performance Management) Storage Administrator (Local Copy) Storage Administrator (Remote Copy) Storage Administrator (Initial Configuration) Storage Administrator (System Resource Management)
Restart data collection tasks	All Resources You can perform actions on tasks that you have created even if All Resources is not allocated.	Admin or Modify For tasks that require Admin permission, you must have Admin permission to register the task.	-
HCS tasks (Tasks & Alerts tab)			
View HCS tasks list	Any	Any	
View HCS tasks related to a storage system, task details, stop tasks, change task schedules, cancel tasks, delete tasks, or move tasks to the history.	All Resources You can perform these actions on tasks that you have created even if All	Admin, Modify, or the user who created the task.	Storage Administrator (System Resource Management)

Function	Resource group	Required roles	
		Admin, Modify, View	Custom (VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, and VSP F1500)
	Resources is not allocated. For VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, and VSP F1500, users can also perform this operation on the default resource group.		
View details of HCS tasks (tasks related to a file server or an HDT monitoring schedule template), change task schedules, stop tasks, cancel tasks, delete task, move tasks to the history.	All Resources	Admin or Modify	-
Restart HCS tasks	All Resources You can perform these actions on tasks that you have created even if All Resources is not allocated. For tasks that require Admin permission, you must have Admin permission to register the task.	Admin or Modify	-
System tasks			
View system tasks Manage system tasks (when in the Tasks & Alerts tab)	Any	Any	
Manage system tasks (when using Device Manager - Storage Navigator)	All Resources Default resource group of the storage system. You can perform these actions on tasks that you have created even if All Resources is not allocated.	Admin or Modify	Storage Administrator (System Resource Management)
Alerts			
View alerts	Resource group to which the resources of the	Any	

Function	Resource group	Required roles	
		Admin, Modify, View	Custom (VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, and VSP F1500)
	target storage system belong.		
Delete alerts	Resource group to which the resources of the target storage system belong.	Admin or Modify	One of the following: Storage Administrator (Provisioning) Storage Administrator (Performance Management) Storage Administrator (Local Copy) Storage Administrator (Remote Copy) Storage Administrator (Initial Configuration) Storage Administrator (System Resource Management)
Search & reports (CSV)			
Edit shared search conditions, or change them to private.	All Resources	Admin	-
Output CSV files containing resource information.	Resource group to which the target resource belongs.	Any	
Volumes			
Create or delete parity groups (HUS100 or AMS2000)	Resource group to which the target resource belongs	Admin or Modify	-
Create parity group (VSP Fx00 models, VSP Gx00 models)	Default resource group of the storage system.	Admin or Modify	Storage Administrator (Provisioning)
Delete parity group (VSP Fx00 models, VSP Gx00 models)	Resource group to which the target resource belongs.	Admin or Modify	Storage Administrator (Provisioning)
Edit parity groups (VSP Fx00 models, VSP Gx00 models)	Resource group to which the target resource belongs.	Admin or Modify	Storage Administrator (Provisioning)
Assign spare drives	Default resource group of the storage system.	Admin or Modify	Storage Administrator (Provisioning)
Create volumes	Resource group to which the target parity group or the DP pool belongs. Also, the resource group to	Admin or Modify	Storage Administrator (Provisioning)

Function	Resource group	Required roles	
		Admin, Modify, View	Custom (VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, and VSP F1500)
	which the unused LDEV ID belongs.		
Edit volume label Allocate volumes, allocate like volumes, define clustered-host storage, or cancel volume allocation. Shred volumes	Resource group to which the target resource belongs	Admin or Modify	Storage Administrator (Provisioning)
Delete volumes	Resource group to which the LDEV ID of the target volume belongs. Also, the resource group to which the parity group or the DP pool belongs.	Admin or Modify	Storage Administrator (Provisioning)
Create LUSE volumes Release LUSE volumes	Resource group to which the LDEV ID of the target volume belongs	Admin or Modify	-
Volumes - global-active device pairs			
View, modify and add when using Set up Replication/ GAD.	All Resources	Admin You must have Replication Manager Admin permissions.	-
Allocate and unallocate global-active device pairs	The resource group to which the target resource belongs.	Admin or Modify	Storage Administrator (Provisioning)
Checking global-active device volume status	All Resources	Admin, Modify, or View You must have Replication Manager View permissions.	-
Delete remote paths	The resource group to which the resources of the target storage system belong.	Admin or Modify	Storage Administrator (Remote Copy)
Delete Quorum disks	The resource group to which the resources of the target storage system belong.	Admin or Modify	Storage Administrator (Provisioning)
Unvirtualize Quorum disks	Resource group that can access the external port and parity group of an internal storage system, and internal volumes that belong to the external parity groups, and host group, volume, target port of an external storage system.	Admin or Modify	Storage Administrator (Provisioning)

Function	Resource group	Required roles	
		Admin, Modify, View	Custom (VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, and VSP F1500)
Unallocate command device volumes Change a command device to a basic volume (delete command device)	The resource group to which the target resource belongs.	Admin or Modify	Storage Administrator (Provisioning)
Volumes and copy pairs - nondisruptive migration			
Create and delete migration projects Edit resource group mapping Edit port mapping Edit external path setting	All Resources	Admin	-
Create, edit, and cancel migration jobs Delete volumes for migration jobs	All Resources	Admin or Modify	-
Perform the following steps for migration jobs or resources in migration jobs: • Virtualize and Allocate Volumes • Prepare Target Secondary Volumes	All Resources	Admin For the Modify role, tasks can be performed if enough resources have already been added to the virtual storage machine.	-
Perform all steps for migration jobs or resources in migration jobs, *except* : • Virtualize and Allocate Volumes • Prepare Target Secondary Volumes	All Resources	Admin or Modify	-
Viewing migration jobs	All Resources	Any	-
External storage systems			
Virtualize volumes of a registered external storage system.	Resource group that can access the external port and LDEV ID of an internal storage system, and the port, LDEV ID, and Host Group ID of an external storage system.	Admin or Modify	Storage Administrator (Provisioning)
Virtualize volumes of an unregistered external storage system.	Resource group that can access external port and LDEV ID of an internal storage system.	Admin or Modify	Storage Administrator (Provisioning)

Function	Resource group	Required roles	
		Admin, Modify, View	Custom (VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, and VSP F1500)
Unvirtualize volumes of a registered storage system.	Resource group that can access the external port and parity group of an internal storage system, and internal volumes that belong to the external parity groups, and host group, volume, target port of an external storage system.	Admin or Modify	Storage Administrator (Provisioning)
Unvirtualize volumes of an unregistered storage system.	Resource group that can access the external port and parity group of an internal storage system and the internal volumes belong to external parity groups.	Admin or Modify	Storage Administrator (Provisioning)
Pools/tiers			
Create or delete pools Edit or shrink pools	Resource group to which the target DP pool belongs.	Admin or Modify	Storage Administrator (Provisioning)
Expand pools	Resource group to which the target DP pool belongs. Also, the resource group to which the following belongs: the LDEV IDs of the DP pool volumes to be added to the DP pool.	Admin or Modify	Storage Administrator (Provisioning)
Create, edit, or delete tiers	Resource group to which the target resource belongs	Admin or Modify	One of the following: Storage Administrator (Provisioning) Storage Administrator (Performance Management) Storage Administrator (Local Copy) Storage Administrator (Remote Copy) Storage Administrator (Initial Configuration) Storage Administrator (System Resource Management)

Function	Resource group	Required roles	
		Admin, Modify, View	Custom (VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, and VSP F1500)
Edit DP pool name	Resource group to which the target resource belongs	Admin or Modify	Storage Administrator (Provisioning)
Expand DP volume size or reclaim zero pages.	Resource group to which the LDEV ID of the DP target volume belongs. Also, the resource group to which the DP pool belongs.	Admin or Modify	Storage Administrator (Provisioning)
Manually start and stop the monitoring and tier relocation of HDT pools.	Resource group to which the target HDT pool belongs.	Admin or Modify	Storage Administrator (Provisioning)
Create, edit, or delete schedule templates for HDT pool monitoring and tier relocation.	All Resources	Admin or Modify	-
Apply schedule template for HDT pool monitoring and tier relocation.	Resource group to which the LDEV ID of the target HDT volume belongs. Also, the resource group to which the HDT pool belongs.	Admin or Modify	Storage Administrator (Provisioning)
Edit tier relocation for HDT volumes.	Resource group to which the target LDEV ID of the HDT volume belongs.	Admin or Modify	Storage Administrator (Provisioning)
Edit tiering policy for HDT volumes.	Resource group to which the target LDEV ID of HDT volumes belong.	Admin or Modify	Storage Administrator (Provisioning)
Change tiering policy definition for HDT volumes (customize tiering policies).	Resource group to which the storage system of HDT volumes, and to which the target tiering policy is applied belongs.	Admin or Modify	Storage Administrator (Provisioning)
Verify whether tiering policies for HDT volumes have been applied.	Resource group to which the LDEV IDs of the volumes to which the tiering policy has been applied belongs.	Any	
Manage data placement profiles for HDT volumes (creating, updating, editing, deleting, applying profiles, releasing applied profiles, and setting schedules).	Resource groups to which the LDEV ID of every HDT volume included in the target data placement profile belongs.	Admin or Modify	Storage Administrator (Provisioning)

Function	Resource group	Required roles	
		Admin, Modify, View	Custom (VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, and VSP F1500)
Edit tier rank for external HDT pool volumes.	Resource group to which the target HDT pool belongs.	Admin or Modify	Storage Administrator (Provisioning)
VVol environment for block storage			
Create storage container	Resource group to which the LDEV IDs of DP pool volumes belong.	Admin or Modify	One of the following: Storage Administrator (Provisioning) Storage Administrator (Performance Management) Storage Administrator (Local Copy) Storage Administrator (Remote Copy) Storage Administrator (Initial Configuration) Storage Administrator (System Resource Management)
Edit/delete storage container	Resource group to which the target storage containers belong.	Admin or Modify	
View storage container	Resource group to which the target storage containers belong.	Any	
Set up protocol endpoint volumes	Resource group to which the target LDEV ID belongs.	Admin or Modify	Storage Administrator (Provisioning)
View protocol endpoint/ VVol	Resource group to which the target resource belongs.	Any	
Settings used by the VASA provider to access HCS	All Resources	Admin or Modify	
VVol environment for file storage			
View the VASA provider for file VVols	All Resources	Any	-
Add, edit, or delete the VASA provider for file VVols	All Resources	Admin or Modify	-
View storage containers for file VVols	All Resources	Any	-
Add, edit, or delete storage containers for file VVols	All Resources	Admin or Modify	-

Function	Resource group	Required roles	
		Admin, Modify, View	Custom (VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, and VSP F1500)
File servers			
Edit file server names Delete file server	Any	Admin or Modify	One of the following: Storage Administrator (Provisioning) Storage Administrator (Performance Management) Storage Administrator (Local Copy) Storage Administrator (Remote Copy) Storage Administrator (Initial Configuration) Storage Administrator (System Resource Management)
File servers - NAS Platform family			
Create, expand, or delete a storage pool Change a storage pool name (NAS Platform v11.3 or later)	All Resources Even if All Resources is not assigned, tasks can be performed by following the same steps used for NAS Platform versions earlier than v11.3. In this case, the SMU permission appropriate to the task is required.	Admin or Modify For the View role, tasks can be performed by following the same steps used for NAS Platform versions earlier than v11.3. In this case, the SMU permission appropriate to the task is required.	-
Create, expand, or delete a storage pool (NAS Platform versions earlier than 11.3)	Any	Any In addition, the SMU permission appropriate to the task is required.	
Add, edit, or delete file shares Expand, mount, unmount, or delete a file system (NAS Platform v11.1 or later)	All Resources Even if All Resources is not assigned, tasks can be performed by following the same steps used for NAS Platform versions earlier than v11.1. In this case, the SMU permission appropriate to the task is required.	Admin or Modify For the View role, tasks can be performed by following the same steps used for NAS Platform versions earlier than v11.1. In this case, the SMU permission appropriate to the task is required.	-
Add, edit, or delete file shares	Any	Any In addition, you must have the SMU permission for the task.	

Function	Resource group	Required roles	
		Admin, Modify, View	Custom (VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, and VSP F1500)
Expand, mount, unmount, or delete a file system (NAS Platform versions earlier than v11.1)			
Set up or modify Admin services EVS (NAS Platform v10.2.3071 or later)	Any	Admin or Modify	One of the following: Storage Administrator (Provisioning) Storage Administrator (Performance Management) Storage Administrator (Local Copy) Storage Administrator (Remote Copy) Storage Administrator (Initial Configuration) Storage Administrator (System Resource Management)
Create a file system (NAS Platform v10.2.3071 or later)	All Resources Even if All Resources is not assigned, tasks can be completed by following the same steps used for NAS Platform versions earlier than v10.2.3071. In this case, the SMU permission appropriate to the task is required.	Admin or Modify For the View role, tasks can be performed by following the same steps used for NAS Platform versions earlier than v10.2.3071. In this case, the SMU permission appropriate to the task is required.	-
Create a file system (NAS Platform versions earlier than v10.2.3071)	Any	Any In addition, you must have the SMU permission to perform these tasks.	
Change settings of a storage pool Edit a file system Edit the system drive Create or edit a system drive group (NAS Platform v10.1.3071 or later)	Any	Any In addition, the SMU/NAS Manager permission appropriate to the task is required.	
File servers - HDI & HNAS F			

Function	Resource group	Required roles	
		Admin, Modify, View	Custom (VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, and VSP F1500)
Create a file system (when linked with File Services Manager).	Resource group to which the target resource belongs	Admin or Modify With any role, you can use volumes allocated to the file server. In addition, you must have File Services Manager Admin permissions.	One of the following: Storage Administrator (Provisioning) Storage Administrator (Performance Management) Storage Administrator (Local Copy) Storage Administrator (Remote Copy) Storage Administrator (Initial Configuration) Storage Administrator (System Resource Management)
Expand or delete a file system (when linked with File Services Manager).	Resource group to which the target resource belongs	Any In addition, you must have File Services Manager Admin permissions.	
Add, edit, or release file shares (when linked with File Services Manager).	Any	Any In addition, you must have File Services Manager Admin permissions.	
Replication			
View, modify, and add resources for setting up a replication configuration.	All Resources	Admin You must have Replication Manager Admin permissions.	-
Create, expand, or shrink a copy topology from the Replication tab	All Resources	Admin or Modify You must have Replication Manager Admin or Modify permissions.	-
Expand and shrink a copy topology by changing the type	All Resources	Admin or Modify You must have Replication Manager Admin or Modify permissions.	-
Change the copy topology name	All Resources	Admin or Modify You must have Replication Manager Admin or Modify permissions.	-
Create, edit, or delete alert settings of copy groups	All Resources	Admin or Modify	-

Function	Resource group	Required roles	
		Admin, Modify, View	Custom (VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, and VSP F1500)
		You must have Replication Manager Admin or Modify permissions.	
Monitor the status of copy topology	All Resources	Admin, Modify, or View You must have Replication Manager View permissions.	-
Refresh configuration information of information sources	All Resources	Admin You must have Replication Manager Admin permissions.	-
Add, edit or delete command devices (no Replication Manager license)	Resource group to which the LDEV ID of the target command device belongs	Admin or Modify	Storage Administrator (Provisioning)
Define copy pairs, or change the status of copy pairs (no Replication Manager license)	Resource group to which the LDEV ID of the target pair volumes and command device allocated to the pair management server belongs	Admin or Modify	For local copy: Storage Administrator (Local Copy) For remote copy: Storage Administrator (Remote Copy)
Update the configuration and performance information of a copy group (UR Performance Analysis window of the Replication tab), set the network bandwidth (UR Performance Analysis window of the Replication tab), set M-JNL threshold values (UR Performance Analysis window of the Replication tab).	All Resources	Admin or Modify You must have Replication Manager View permissions.	-
Confirm and analyze Universal Replicator performance (UR Performance Analysis window of the Replication tab), export the analysis results (UR Performance Analysis window of the Replication tab).	All Resources	Admin, Modify, or View You must have Replication Manager View permissions.	-
Mobility (migration)			
View the View Tier Properties dialog box for an	Resource group to which the target HDT pool belongs.	Any	

Function	Resource group	Required roles	
		Admin, Modify, View	Custom (VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, and VSP F1500)
HDT pool or each tiering policy of HDT volumes.			
View the View Tier Properties dialog box for an HDT volume.	Resource group to which the target LDEV ID of an HDT volume belongs.	Any	
View the Trend Chart and output the CSV file from the Mobility tab.	Resource group to which the target resource belongs	Any	
Data migration	Resource group to which the target resource belongs.	Admin or Modify	Storage Administrator (Provisioning)
Virtual ID			
Display Virtual ID information	Resource group to which the target resource belongs	Any	
Resources of virtual storage machines			
View information of virtual storage machines	Default resource group of the storage system.	Any	
Create, edit or delete virtual storage machines. Move resources in a virtual storage machine to a default virtual storage machine, or resources in the default virtual storage machine to another virtual storage machine.	All Resources	Admin	-
Manage resources (including adding or deleting virtual information of LDEVs) in a virtual storage machine (except the default virtual storage machine).	Resource group to which the target resource belongs in the virtual storage machine.	Admin or Modify	Storage Administrator (Provisioning)
Analytics (Analytics tab)			
Identify the existence of performance problems of storage systems, analyze MP Blades/Units performance, analyze storage performance on a host basis, perform Health Check immediately, and view and export Health Check reports.	All Resources	Admin, Modify, or View You must have Tuning Manager View permissions.	-

Function	Resource group	Required roles	
		Admin, Modify, View	Custom (VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, and VSP F1500)
Execute Health Checks on a schedule (Analytics tab), change the threshold value for performance analysis of storage systems (Analytics tab).			
Delete a Health Check report	All Resources	Admin, Modify, or View In addition, you must be logged in with a user ID that has Tuning Manager Admin permissions or with the user ID that generated the Health Check report.	-

For HCS, custom roles (Storage, Security, and Audit Log roles, as well as the Support Personnel and User Maintenance roles for the VSP Fx00 models, VSP Gx00 models,) provide specific permissions for performing specific tasks (functions) on VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, and VSP F1500. Additional tasks (functions) are available when selecting System GUI from menus or application panes to open Device Manager - Storage Navigator, and are listed in the table below. Tasks related to system administration and maintenance are available for the VSP Fx00 models and VSP Gx00 models by launching the maintenance utility from the list of storage systems on the Resources tab.

Table 8 Required roles or custom roles for performing specific functions

Functions	Resource group	Required roles	
		Admin, Modify	Custom (VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, and VSP F1500)
Provisioning related operations: • Configuring caches • Configuring LDEVs, pools, and virtual volumes • Formatting and shredding LDEVs • Configuring external volumes • Creating and deleting quorum disks used in a global-active device environment. • Creating and deleting global-active device pairs. • Configuring alias volumes for Compatible PAV	Resource group to which the resources of the target system belong.	Admin or Modify	Storage Administrator (Provisioning)

Functions	Resource group	Required roles	
		Admin, Modify	Custom (VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, and VSP F1500)
• Configuring Dynamic Provisioning • Configuring host groups, paths, and WWNs/ iSCSI • Configuring Volume Migration except splitting Volume Migration pairs when using CCI. • Configuring access attributes for LDEVs • Configuring LUN security			
Performance related operations: • Configuring monitoring • Starting and stopping monitoring	Resource group to which the resources of the target system belong.	Admin or Modify	Storage Administrator (Performance Management)
Local copy management: • Performing pair operations for local copy • Configuring environmental settings for local copy • Splitting Volume Migration pairs when using CCI	Resource group to which the resources of the target system belong.	Admin or Modify	Storage Administrator (Local Copy)
Remote copy management: • Remote copy operations in general • Managing global-active device pairs (except for creation and deletion).	Resource group to which the resources of the target system belong.	Admin or Modify	Storage Administrator (Remote Copy)
Initial configuration: • Configuring settings for storage systems • Configuring settings for SNMP • Configuring settings for e-mail notification • Configuring settings for license keys • Viewing, deleting, and downloading storage configuration reports • Acquiring all the information about the storage system and refreshing	Default resource group of the storage system.	Admin or Modify	Storage Administrator (Initial Configuration)
System resource management: • Configuring settings for CLPR • Configuring settings for MP Blade • Deleting tasks and releasing exclusive locks of resources • Completing SIMs • Configuring attributes for ports • Configuring LUN security • Configuring Server Priority Manager • Configuring tiering policies	Default resource group of the storage system.	Admin or Modify	Storage Administrator (System Resource Management)
For global-active device: • Setting the reserved attribute for a volume to be used in a global-active device pair.	Default resource group of the storage system.	Admin	Security Administrator (View & Modify)

Functions	Resource group	Required roles	
		Admin, Modify	Custom (VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, and VSP F1500)
For encryption key and authentication management (does not allow user management): • Creating an encryption key, configuring encryption • Viewing and switching the location at which to create an encryption key • Backing up and restoring an encryption key • Deleting an encryption key that is backed up on the key management server, or key service processor. • Viewing and changing the password policies for backing up an encryption key on the management client. • Configuring the certificate used for SSL communication on the management client*. • Configuring the Fibre Channel authentication (FC-SP)			
Viewing storage encryption key and authentication: • Viewing information about encryption settings • Viewing information about encryption keys on the key management server, or key service processor.	Default resource group of the storage system.	Admin	Security Administrator (View & Modify) or Security Administrator (View Only)
Storage system audit log management: • Configuring audit log settings • Downloading audit logs	Default resource group of the storage system.	Admin	Audit Log Administrator (View & Modify)
Storage system audit log access: • Viewing storage audit log settings • Downloading audit logs	Default resource group of the storage system.	Admin	Audit Log Administrator (View & Modify) or Audit Log Administrator (View Only)
• Add or remove hardware • View hardware configuration information • Perform maintenance operations by using a maintenance PC (not including operations for the maintenance utility) • Replace parts (including basic maintenance of HDDs) • View SIM/SSB • Install OS security patches • Update OSes • Restore MP • Turn power supply on or off • Manage storage systems (View)	Default resource group of the storage system.	Admin	Support Personnel These functions are available for the exclusive use of service representatives.
• Add or remove hardware • Install OS security patches	Default resource group of the storage system.	Admin	User Maintenance

Functions	Resource group	Required roles	
		Admin, Modify	Custom (VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, and VSP F1500)
• Update OSes • Replace parts • Perform basic maintenance • Restore MP • View SIM and SSB • Turn power supply on or off • Manage storage systems (View)			These functions are available for the exclusive use of service representatives.
* For VSP G1000, VSP G1500 and VSP F1500, when a user account for logging in to the SVP or Command Control Interface (CCI) is authenticated by HCS, if the user account created in HCS is assigned the Security Administrator (View & Modify) role, that user account can be used to open the Tool Panel and configure the certificate. For details about this procedure, see the *System Administrator Guide* for your storage system.			

Related concepts

- About user accounts and controlling access to resources on page 88

Related tasks

- Changing permissions for a user account on page 95
- Editing the profile for a user account on page 93

Related references

- User group roles on page 110
- Custom roles on page 111

Creating user groups

You can create user groups and assign resource groups either immediately, or later as needed. The assigned resource group rights provide user group members with view, modify, admin, or custom roles for displaying, modifying, or managing the storage resources for this group.

Creating a user group also creates a private logical group for that user group.

Procedure

1. On the **Administration** tab, select **User Groups**.
   Existing user groups are displayed.
2. Click **Create User Group** and enter a name and description.
3. Click **Add Users** to select the user group members.
4. Click **Add Resource Groups** to assign specific storage system resources to the user group, and for each resource group set the role (permissions) to one of the following:
   - View

- Modify
- Admin
- Custom

5. Click **OK**.
   You can click the user group name to verify group membership (list of users) on the **Users** tab. On the **Resource Groups** tab, you can verify the resource groups and roles specified during user group creation. **Add Resource Groups**, **Remove Resource Groups**, and **Edit Roles** buttons allow you to modify resources and access rights for the user group.

Result

The new user group is displayed and a new private logical group is created.

Related concepts

- About user groups on page 107

Related tasks

- Editing a user group on page 135
- Assigning resource groups and roles to a user group on page 136
- Deleting user groups on page 136

Editing a user group

As user information and membership in user groups change, you can update this information.

Procedure

1. On the **Administration** tab, select **User Groups**.
2. Select the target user group and click **Edit User Group**.

   Note: If the **Edit User Group** button is not activated, the selected user group cannot be edited. This occurs when selecting a built in user group for which editing is not allowed.

3. Change the name of the user group, or the user who manages the user group.

Result

Editing changes appear in the list in the User Groups tab.

Related concepts

- About user groups on page 107
- About access control on page 99

Related tasks

- Assigning resource groups and roles to a user group on page 136

Assigning resource groups and roles to a user group

You can assign resource groups and roles to a user group.

Procedure

1. Select the **Administration** tab and click on **User Groups**.
2. In the **User Groups** tab, click the name of the target user group.
3. In the **Resource Groups** tab, click **Add Resource Groups**.
4. Set the roles to be assigned to the user group. The assigned resource groups and roles are displayed in detail by clicking the link for the user group name.

Related tasks

- Creating resource groups on page 104
- Creating user groups on page 134

Changing a user's user group

You can modify information about the user groups that a user belongs to.

Procedure

1. On the **Administration** tab, select **User Groups**.
2. In the **Users** tab select the target user group and click **Assign User Groups**.
3. Change the user groups to which a user belongs.

Result

Changes will appear in the Users tab.

Related concepts

- About access control on page 99

Related tasks

- Creating user groups on page 134
- Assigning resource groups and roles to a user group on page 136
- Editing a user group on page 135

Deleting user groups

Deleting a user group deletes the associated private logical group and all logical groups that belong to that private logical group. Deleting a user group does not delete the related resource groups.

Procedure

1. On the **Administration** tab, select **User Groups**.
2. Select the user groups to delete.
3. Click **Delete Groups**.

> **Note:** If the **Delete Groups** button is not activated, the selected user group cannot be deleted. This occurs when selecting a built in user group that is not allowed to be deleted, or if the user group is associated with another group, such as a resource group, and deletion is not allowed.

4. Click **OK**.

Result

The user groups you deleted no longer appear in the list of user groups.

5

Provisioning storage

This module describes provisioning storage.

- Creating a storage operating environment
- Allocating storage
- Creating and deleting parity groups
- Creating and deleting volumes
- Virtualizing external storage
- Virtualizing storage capacity (HDP/HDT)
- Virtualizing storage tiers (HDT)
- Allocating and unallocating volumes
- Configuring Fibre Channel and iSCSI ports
- Managing LUN paths
- Managing NAS Platform file systems and shares
- Managing NAS Platform F and Data Ingestor file systems and shares
- Managing copy pairs
- Managing command devices

Creating a storage operating environment

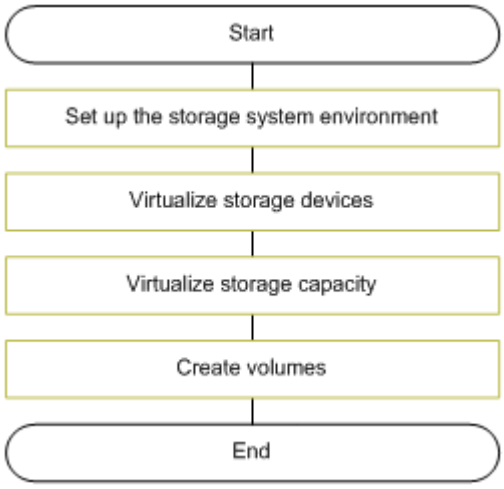
To be able to use volumes of a storage system from hosts, file servers, or NAS modules configure the operating environment beforehand. The workflow for configuring the environment differs depending on the scale and operation method of the storage systems.

For example, you can create parity groups for supported mid-range storage, then create and allocate volumes. When the volumes are allocated, file systems can be created and mounted. For enterprise storage, the parity groups will already exist as a result of installation. You can allocate volumes for registered hosts, file servers, and NAS modules.

In some environments, you can virtualize storage devices so that multiple, different storage systems can be used as a single storage system.

If your storage system supports DP pools, you can create DP pools such that virtual volumes can be allocated to hosts. By virtualizing storage devices and storage capacity, you can decrease both management and operational costs by more effectively using physical resources.

The following figure shows an example workflow for building an environment for performing operations with virtualized storage devices or volumes, when using VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, VSP F1500, VSP, and HUS VM.



Related concepts

- About creating parity groups on page 142
- About creating volumes on page 149
- About virtualizing and unvirtualizing volumes on page 160
- About virtualizing storage capacity on page 171
- About virtualizing storage tiers on page 187
- Allocating storage on page 141

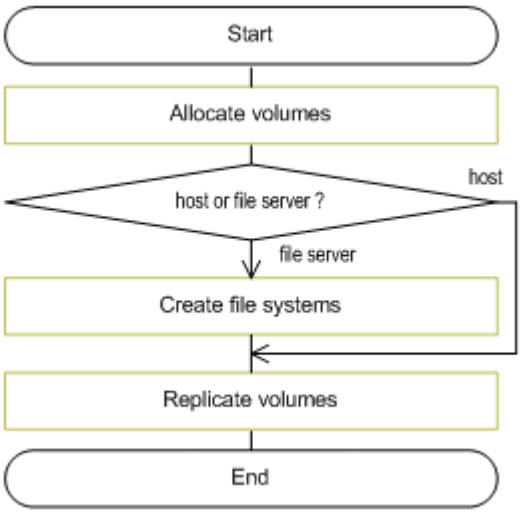
Allocating storage

Volumes are allocated to hosts for applications needing storage.

A variety of methods for allocating volumes is available. For example, you can select one or more hosts, then identify and allocate existing volumes. You can also select one or more volumes, then identify the host that needs the volumes. You can also establish I/O paths between hosts and volumes when you allocate volumes. For file servers or NAS modules, allocate volumes, and then create a file system or a file share.

To prepare for loss of data caused by disk failure, disasters, or other issues, you can manage the redundancy of important operational data by creating a replication environment for volumes within a storage system or between storage systems, as necessary.

Note: Replication services apply to host and file server volumes, but not file server file systems.



Related concepts

- About allocating volumes on page 201
- About managing NAS Platform file systems and shares on page 246
- About managing NAS Platform F and Data Ingestor file systems and shares on page 273
- About replicating volumes on page 288
- Creating a storage operating environment on page 140

Creating and deleting parity groups

This module describes how to create and delete parity groups on VSP Fx00 models, VSP Gx00 models, Hitachi Unified Storage (HUS) 100 and Adaptable Modular Storage (AMS) 2000 family of storage systems.

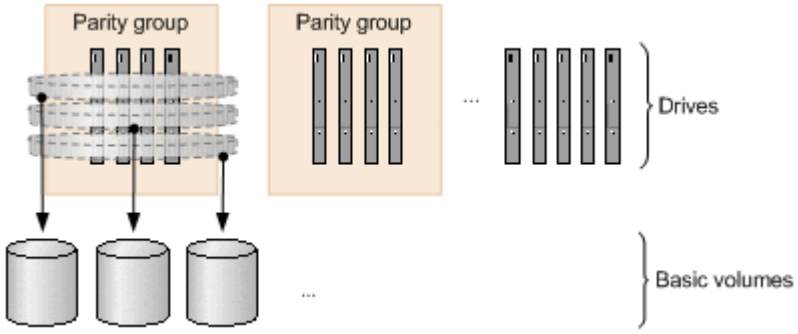
About creating parity groups

Creating parity groups is necessary for specific mid-range storage requirements. Hitachi Command Suite supports parity group management only for VSP Fx00 models, VSP Gx00 models, and Hitachi Unified Storage (HUS) 100 and Adaptable Modular Storage (AMS) 2000 family of storage systems.

You must first create parity groups before you can create and then allocate the basic volumes of a storage system to a host. A parity group is created from multiple drives. For example, create parity groups in the following situations:

- When introducing a new storage system with installed drives.
- When allocating volumes to a host or creating a volume and the displayed list of parity groups indicates that there is not enough capacity, and drives are added.
- When creating a parity group for use in a NAS Platform storage pool.

The following graphic illustrates the relationship between parity groups, volumes, and drives:



Related concepts

- About deleting parity groups on page 145
- Creating a storage operating environment on page 140

Related tasks

- Creating parity groups on page 143

Related references

- Create parity groups dialog box on page 144

- Conditions for parity groups that are used in storage pools on page 253

Creating parity groups

You can create and manage parity groups for VSP Fx00 models, VSP Gx00 models storage systems, and for the Unified Storage 100 and Adaptable Modular Storage 2000 family of storage systems.

Before you begin

- Register the target storage system
- Identify the target storage system name
- Identify the required number of parity groups, and related capacities
- Identify desired RAID levels, drive types and speeds (RPM)

Procedure

1. On the **Resources** tab, right-click on a supported storage system and choose **Create Parity Groups**.
2. Depending on your storage system type, do one of the following:
   - In the **Create Parity Groups** dialog box, select drive types, RAID level, and parity group options. Proceed to step 3.
   - In the **Create Parity Groups** wizard page, select drive types, RAID level, and parity group options, click **Next**, verify the settings, and then click **Finish** to register a System task. Proceed to step 7.
3. Click **Show Plan** and confirm that the information in the plan summary is correct. If changes are required, click **Back**.
4. (Optional) Update the task name and provide a description.
5. (Optional) Expand **Schedule** to specify the task schedule.

   You can schedule the task to run immediately or later. The default setting is **Now**. If the task is scheduled to run immediately, you can select **View task status** to monitor the task after it is submitted.
6. Click **Submit.**
   If the task is scheduled to run immediately, the process begins.
7. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.

Result

Created parity groups are added to the target storage system parity group list.

Related concepts

- About creating parity groups on page 142
- About deleting parity groups on page 145

Related references

- Create parity groups dialog box on page 144

Create parity groups dialog box

Creating parity groups is necessary for storage systems such as Unified Storage 100 and Adaptable Modular Storage 2000 family of storage systems that have been configured with additional physical disk capacity, but still require parity group RAID configuration before creating and allocating volumes.

When you enter the minimum required information in this dialog box, the Show Plan button activates to allow you to review the plan. Click the Back button to modify the plan to meet your requirements.

The following table describes the dialog box fields, subfields, and field groups. A *field group* is a collection of fields that are related to a specific action or configuration. You can minimize and expand field groups by clicking the double-arrow symbol (>>).

As you enter information in a dialog box, if the information is incorrect, errors that include a description of the problem appear at the top of the box.

Table 9 Create parity groups dialog box

Field	Subfield	Description
Storage System	-	This field will either display the selected storage system name, or prompt the user to select the storage system from a list.
Available Drive Types	-	Available drives are displayed, including drive type, drive speed, form factor, the number of available drives (which influences configurable RAID levels for a given set of drives), and total capacity.
RAID Level	RAID	Valid RAID levels are displayed for the selected drive type and count. For example, RAID 0, 1, 1+0, 5 and 6 are valid RAID configurations when 6 drives are available (for example, HUS 100). With 2 drives, only RAID 0 and 1 would be valid. These changes are reflected in the dialog box based on the selected storage system (and supported RAID levels), and selected drive type and count.
	Drive	Additionally, valid drive combinations are displayed for a given RAID level. For example, with 6 available drives, RAID 5 can be configured as 2D+1P, 3D+1P, 4D+1P, or 5D+1P (for example, HUS 100). Using 2D+1P an administrator could create two parity groups instead of one, and use them for different requirements, or the administrator could create a single large 5D+1P parity group using all available drives.
Parity Group capacity	-	This displays the usable capacity of the selected RAID level and drive configuration, and the size of the physical drives.
No. of Parity Groups	-	Specify the number of parity groups to create. The maximum is determined by the selected RAID configuration and the total number of available drives.
>> Advanced Options	Parity Group Number:	Auto parity group number assignments are made by the system, and start at 1 by default.

Field	Subfield	Description
	Auto Manual	Manual parity group number assignments are made by the administrator.
	Drive Selection: Automatic Manual	Automatic drive selection means the drives needed for the parity group will be chosen by the system. Manual drive selection allows the administrator to determine which physical drives will be used for a parity group. For example, for a RAID 5, 2D+1P parity group, the system will typically select the first three (3) drives seen. Clicking Select Drives allows the administrator to list all drives, and select the desired drives for the parity group.
	Selected Drives	This panel displays the automatically or manually selected drives for parity group creation, as discussed above. The display includes drive number, tray number, HDU number, drive model and related firmware.

Related concepts

- About creating parity groups on page 142

Related tasks

- Creating parity groups on page 143

About deleting parity groups

Deleting parity groups is one of the parity group management tasks you can perform. Hitachi Command Suite supports parity group management only for VSP Fx00 models, VSP Gx00 models, Unified Storage 100 and Adaptable Modular Storage 2000 family of storage systems.

Free drives are necessary to create and use DP pools in VSP Fx00 models, VSP Gx00 models, HUS 100 and AMS 2000 family of storage systems. If there are not enough free drives when creating a DP pool, users can delete unnecessary parity groups to increase the number of free drives.

Caution: When a parity group is deleted, the data in the parity group volumes is deleted.

Related tasks

- Deleting parity groups on page 145

Deleting parity groups

You can delete parity groups from VSP Fx00 models, VSP Gx00 models, Unified Storage 100 and Adaptable Modular Storage 2000 family of storage systems.

Note: When you remove an encrypted parity group in HUS 150 systems, this cancels the encryption for all drives that make up the target parity group, and releases the encryption on all volumes in the parity group.

Before you begin

- Identify the target storage system that is associated with the parity groups you want to delete.
- Identify the target parity groups to delete.
- Unallocate all volumes in the target parity groups.

Procedure

1. On the **Resources** tab, select **Storage Systems**.
2. Expand the tree, select the target storage system, and select **Parity Groups**.
3. From the list of parity groups, select one or more target parity groups to delete, and click **Delete Parity Groups**.

   Tip: When deleting parity groups that compose a concatenated parity group, select the parity group that has the smallest number within the concatenated parity group. When you do this, all concatenated parity groups are deleted.

4. Depending on your storage system type, do one of the following:
   - In the **Delete Parity Groups** dialog box, confirm that the information displayed is correct. Optionally, update the task name and provide a description.
   - In the **Delete Parity Groups** wizard page, confirm that the information displayed is correct, and then click **Apply**.
5. Expand **Schedule** to specify the task schedule.
   The task can be run immediately or scheduled for later. The default setting is **Now**.
6. Click **Submit**.
   If you selected **Now**, the delete parity groups process begins.
7. You can check the progress and the result of the delete parity groups task on the **Tasks & Alerts** tab.
   Verify the results for each task by viewing the details of the task.

Result

The deleted parity group no longer appears in the parity groups list for the target storage system.

Related concepts

- About deleting parity groups on page 145

Spare drives and copy-back mode

Spare drives are mounted in Hitachi Virtual Storage Platform G200, G400, G600, G800 or Hitachi Virtual Storage Platform F400, F600, F800 storage systems separately from the data drives for usual write and read operation. If

a failure occurs in a data drive, data is copied to a spare drive so that the storage system can continue to be available.

Copy-back mode determines whether data on a spare drive is returned to a data drive after recovery. If copy-back mode is enabled, when a failed drive is recovered, data is moved from the spare drive to the recovered drive. If copy-back mode is disabled, when a failed drive is recovered, data is left on the spare drive. This mode can be specified for each parity group.

The following maximum number of spare drives can be specified for each storage system:

- VSP G200: 16
- VSP G400, G600 : 32
- VSP G800: 64
- VSP F400, F600, F800: 8

Related tasks

- Assigning spare drives on page 147
- Releasing spare drives on page 148
- Editing copy-back mode on page 148
- Enabling accelerated compression on page 149

Assigning spare drives

Spare drives are used in Hitachi Virtual Storage Platform G200, G400, G600, G800 or Hitachi Virtual Storage Platform F400, F600, F800 storage systems to store data if a data drive fails.

Procedure

1. On the **Resources** tab, click **Storage Systems**, and then expand **All Storage Systems** and the target storage system.
2. Right click **Parity Groups**, and then select **Assign Spare Drives**.
3. In the **Available Drives** table, select one or more drive boxes, and then click **Add**.

   The selected drive boxes are moved to the **Selected Spare Drives** table.
4. Click **Finish**.

   The **Confirm** window opens.
5. In the **Confirm** window verify the settings.
6. In the **Task Name** field enter a unique name or accept the default, and then click **Apply**.

   If **Go to tasks window for status** is checked the **Tasks** window opens.

Related tasks

- Releasing spare drives on page 148

Releasing spare drives

You can release spare drives in Hitachi Virtual Storage Platform G200, G400, G600, G800 or Hitachi Virtual Storage Platform F400, F600, F800 storage systems so they can be used as data drives.

Procedure

1. On the **Resources** tab, click **Storage Systems**, and then expand **All Storage Systems** and the target storage system.
2. Right click **Parity Groups**, and then select **Assign Spare Drives**.
3. In the **Selected Spare Drives** table, select one or more drive boxes, and then click **Remove**.

   The selected drive boxes are moved to the **Available Drives** table.
4. Click **Finish**.

   The **Confirm** window opens.
5. In the **Confirm** window verify the settings.
6. In the **Task Name** field enter a unique name or accept the default, and then click **Apply**.

   If **Go to tasks window for status** is checked the **Tasks** window opens.

Related tasks

- Assigning spare drives on page 147

Editing copy-back mode

Copy-back mode determines whether data on spare drives in VSP Gx00 models or VSP Fx00 models storage systems is restored after a failed drive is recovered.

When editing the copy-back mode of parity groups that are part of a concatenated parity group, select the parity group that has the smallest number. When you do this, the copy-back modes of all parity groups in the concatenated parity group are changed.

Procedure

1. On the **Resources** tab, click **Storage Systems**, and then expand **All Storage Systems** and the target storage system.
2. Click **Parity Groups**.
3. Select one or more parity groups, and then click **Edit Parity Groups**.
4. In the **Edit Parity Groups** window, modify the applicable copy-back mode settings.
5. Click **Finish**.
6. Verify the setting in the **Confirm** window.
7. In the **Task Name** field, enter a unique name (or use the default), and then click **Apply**.

If the **Go to tasks window for status** is checked, then the **Tasks** window opens.

Enabling accelerated compression

Accelerated compression can be set for VSP Gx00 models, VSP Fx00 models, VSP G1000, VSP G1500, and VSP F1500 storage systems.

You can enable accelerated compression when you modify parity groups. If you enable accelerated compression on the parity groups that use FMC, you can take advantage of the increased free capacity because data is compressed.

When editing the accelerated compression of parity groups that are part of a concatenated parity group, select the parity group that has the smallest number. When you do this, the accelerated compression of all parity groups in the concatenated parity group are changed.

For more information about accelerated compression, see the *Provisioning Guide* for your storage system.

Procedure

1. On the **Resources** tab, click **Storage Systems**, and then expand **All Storage Systems** and the target storage system.
2. Click **Parity Groups**.
3. Select one or more parity groups, and then click **Edit Parity Groups**.
4. In the **Edit Parity Groups** window, select **Enable**.
5. Click **Finish**.
6. Verify the setting in the **Confirm** window.
7. In the **Task Name** field, enter a unique name (or use the default), and then click **Apply**.
   If the **Go to tasks window for status** is checked, then the **Tasks** window opens.

Creating and deleting volumes

This module describes how to create volumes and delete unallocated volumes.

About creating volumes

You create volumes, then allocate them to a host.

You create volumes by using the available space in a DP pool or parity group. You can then access the volumes when you are ready to allocate them to a host. If, while allocating volumes to a host, no volumes match the specified requirements, volumes are automatically created using the available space. Note that when a basic volume is created, the volume is also formatted at the same time.

Newly created volumes are included in the list of Open-Unallocated volumes until you allocate them to a host.

Because creating volumes takes time, you should create volumes in advance.

Tip: For VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, or VSP F1500 storage systems, you can block volumes separated from parity groups, recover parity groups from errors, and format volumes by using the windows available by clicking the System GUI link. To access the System GUI link, on the Resources tab, right-click Parity Groups for the target storage system, and then select System GUI from the menu. Or, click Parity Groups for the target storage system, and then click the System GUI link that appears in the application pane.

Additionally, you can format, block, and restore volumes, configure command devices, edit command devices, assign MP blades, and force delete copy pairs (TC, UR, and GAD) by using the windows available by clicking the System GUI link. To access these windows, on the Resources tab, right-click Volumes for the target storage system, and then select System GUI from the menu.

You can allocate the deduplication system data volume (DSD volume) to a DP pool or create a DP volume for which the capacity saving function (dedupe and compression) is enabled, by clicking the **System GUI** link: On the **Resources** tab, right-click **DP Pools** for the target storage system, and then select **System GUI** from the menu.

For information about how to perform tasks that are available by clicking System GUI, see the Help for the appropriate window, or the provisioning guide for your system.

When you are linking with Data Ingestor v11.3 or later and volumes are created for creating or expanding storage pools, it is recommended that you create volumes using the Create Storage Pool or Expand Storage Pool dialog boxes. Device Manager can automatically specify the number of volumes and capacity, and create volumes following the best practices for configuring storage pools.

Related concepts

- About deleting unallocated volumes on page 156
- Creating a storage operating environment on page 140

Related tasks

- Creating volumes on page 151

Related references

- Notes on performing quick formats on page 151
- Create Volumes dialog box on page 152

Notes on performing quick formats

A quick format might impose a heavy workload on some components and lower I/O performance of all hosts running in the target storage system.

We recommend running a quick format when system activity is low and major system operations are not running.

We also recommend running a quick format on a maximum of eight volumes at first, and then confirming that the quick format has not lowered host I/O performance. After that, when you perform a quick format on other volumes, we recommend increasing the number of volumes to be formatted in increments of four.

In particular, if the storage system components are configured as follows, the host I/O performance is likely to be lowered when a quick format is performed:

- Components such as cache memory, CHAs (channel adapters), and DKAs (disk adapters) are in the minimum configuration.
- The number of installed components is extremely different among DKCs (controller chassis) or modules within a single storage system.

In these configurations, run a quick format on only one volume at first, review the host I/O performance, and then continue to run a quick format on other volumes one by one.

Related concepts

- About creating volumes on page 149

Related tasks

- Creating volumes on page 151
- Creating HNAS F or HDI file systems (HFSM v3.2 or later) on page 276

Related references

- Create Volumes dialog box on page 152
- Allocate Volumes dialog box on page 213

Creating volumes

For registered storage systems, volumes are created so they can be allocated to hosts.

Before you begin

- Identify the storage system
- Identify the number of volumes to create
- Identify volume types and capacities
- Create and format parity groups and create DP pools.

Procedure

1. On the **Resources** tab you can create volumes from several locations:
   - From **General Tasks**, select **Create Volumes**.
   - Select the storage system, click **Actions**, and select **Create Volumes**.
   - Select the storage system, list existing parity groups, and click **Create Volumes**.
   - Select the storage system, list existing DP pools, and click the **Create Volumes** button or select **Create Volumes** from **Actions**.
2. In the create volumes dialog box, configure volumes and their characteristics.
3. Click **Show Plan** and confirm that the information in the plan summary is correct. If changes are required, click **Back**.
4. (Optional) Update the task name and provide a description.
5. (Optional) Expand **Schedule** to specify the task schedule.

   You can schedule the task to run immediately or later. The default setting is **Now**. If the task is scheduled to run immediately, you can select **View task status** to monitor the task after it is submitted.
6. Click **Submit.**
   If the task is scheduled to run immediately, the process begins.
7. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.

Result

Created volumes are added to the target storage system Open-Unallocated volume list.

Tip: If a task to create multiple volumes fails, some volumes might have been created even when the Status in Volume Information displayed in the Task Details window is other than Completed. Refresh the target storage system information, and then in the Open-Unallocated volume list, confirm whether volumes were created.

Related concepts

- About creating volumes on page 149

Related references

- Notes on performing quick formats on page 151
- Create Volumes dialog box on page 152

Create Volumes dialog box

Newly created volumes are placed in the Open-Unallocated folder of the user selected storage system until they can be allocated to hosts as needed.

When you enter the minimum required information in this dialog box, the Show Plan button activates to allow you to review the plan. Click the Back button to modify the plan to meet your requirements.

The following table describes the dialog box fields, subfields, and field groups. A *field group* is a collection of fields that are related to a specific action or configuration. You can minimize and expand field groups by clicking the double-arrow symbol (>>).

As you enter information in a dialog box, if the information is incorrect, errors that include a description of the problem appear at the top of the box.

Table 10 Create volumes dialog box

Field	Subfield	Description
No. of Volumes	-	Manually enter the number of volumes to create, or use the arrows (click, or click and hold) to increment or decrement the volume count.
Volume Capacity	-	This number (in blocks, MB, GB, or TB) is the capacity to allocate for each volume. The total capacity to be allocated is calculated as No. of Volumes * Volume Capacity and is displayed. If the drive type is FMC and if accelerated compression is enabled for the parity group, then make sure you specify the amount before compression for the volume capacity.
Storage System	-	This field will either display the selected storage system name, or prompt the user to select the storage system from a list.
Volume Type	-	Select the volume type to create. For example Basic Volume, Dynamic Provisioning or Dynamic Tiering. The displayed volume types are determined by your selected storage system. If you do not see an expected volume type, check that you have selected the correct storage system.
Internal/External	-	When volume type is Basic Volume, or Dynamic Provisioning, volumes can be created using available capacity from the selected storage system (internal) or from an external storage system physically connected to the selected storage system (external).
Pool	-	When volume type is Dynamic Tiering, volumes can be created using Select Pool.
>> Advanced Options	Volume Selection	This is displayed if a storage system was selected and the volume type is Basic. Specify whether to use parity groups or free space to create a volume. This is only displayed when using VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, VSP F1500 Virtual Storage Platform, or Unified Storage VM.
	Drive Type	If multiple drive types are displayed, you can designate a specific drive type.
	Drive Speed (RPM)	If multiple drive speeds are displayed, you can designate a specific drive speed, or accept the default of any available speed.

Field	Subfield	Description
	Chip Type	If multiple chip types are displayed, you can designate a specific chip type.
	RAID Level	If multiple RAID levels are displayed, you can designate a specific RAID level, or accept the default of any available RAID level.
	Select Free Space	After selecting a storage system, specifying Basic for the volume type, and Free Space with volume selection, you can specify free space for parity groups when creating volumes.
	Parity Group	When volume type is Basic Volume, based on drive type, drive speed, chip type, and RAID level selections an appropriate parity group is selected and displayed for you. You can also manually select a parity group by clicking Select Parity Group. In the displayed list of parity groups, you can use sort and filter features on columns such as RAID level, or unallocated capacity (or other fields) to identify the preferred parity groups.
	Pool	When volume type is Dynamic Provisioning, volumes can be created using Select Pool. The listed pools can vary depending on drive type, drive speed, chip type, and RAID level selections.
	Tiering Policy Setting	Displays only if Dynamic Tiering is selected as the volume type, and an HDT pool has been selected with Select Pool (see previous Volume Selection section). You can select a specific tier policy for the volume to be allocated, or select All.
	New Page Assignment Tier	For VSP Gx00 models, VSP G1000, VSP G1500, VSP F1500, VSP and HUS VM, selecting this option specifies to which hardware tier the new page of an HDT volume is to be assigned with a specified priority. Within the hardware tiers for which the tiering policy is set, specify *High* for an upper-level hardware tier, *Middle* for a medium-level hardware tier, and *Low* for a low-level hardware tier.
	Relocation Priority	For VSP Gx00 models, VSP G1000, VSP G1500, VSP F1500, VSP and HUS VM, selecting this option specifies whether you want to prioritize the relocation of the data in HDT volumes.
	Full Allocation	VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, VSP F1500, selecting this option allows you to reserve pages that correspond to the specified capacity when you create volumes.
	Label	Volume labels are searchable, and therefore recommended as a way to find volumes. The Initial value is not required, but can be useful for differentiation when you are creating multiple volumes. Reflect a label to the storage system is checked by default so that naming is consistent between HCS and the storage system itself.
	LDEV ID	An LDEV ID can be assigned automatically or manually.
	Format Type	You can request a quick format, or a basic format. Note that during a quick format, the load might become concentrated on some components, lowering the I/O

Field	Subfield	Description
		performance of all hosts that are running in the target storage system.

Related concepts

- About creating volumes on page 149

Related tasks

- Creating volumes on page 151

Related references

- Notes on performing quick formats on page 151

About shredding volume data

Before deleting a volume that you no longer need, completely remove the data from the volume to avoid unauthorized use of information. The data can be removed by shredding or reformatting the volume.

Volume data is shredded by overwriting it repeatedly with dummy data, which securely destroys the original data. Some volumes, such as basic volumes and DP volumes that are allocated to hosts or used for replication, cannot be shredded.

Some storage systems do not support the shredding functionality. For those storage systems, delete volume information by reformatting the volumes.

Related tasks

- Shredding volume data on page 155

Shredding volume data

Specify one or more volumes that are not allocated to a host and shred the data on the volumes.

Caution: You cannot restore shredded data.

Note: Shred during off hours, such as overnight, so that the shredding process does not adversely affect system performance. To verify the standard required times for shredding, see the *Hitachi Volume Shredder User Guide*.

Before you begin

- Identify the storage system name that includes the volumes that you want to shred.
- Identify the volumes to be shredded.

- Unallocate the volumes to be shredded from the host.

Procedure

1. On the **Resources** tab, select **Storage Systems**.
2. Expand the Storage Systems tree, select a storage system, and from the volumes list, select the volumes whose data you want to shred.
3. Click **Shred Volumes**.
4. In the **Shred Volumes** dialog box, check the target volume to be shredded and edit the writing data pattern, if needed.

   If the storage system does not support shredding, the data will be formatted.
5. (Optional) Update the task name and provide a description.
6. (Optional) Expand **Schedule** to specify the task schedule. You can specify the task to run immediately or later. The default setting is **Now**. If the task is scheduled to run immediately, you can select **View task status** to monitor the task after it is submitted.
7. Click **Submit**. If the task is scheduled to run immediately, the process begins.
8. You can check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.

Result

When the task completes, the data is shredded or reformatted from the volume.

Related concepts

- About shredding volume data on page 155

Related tasks

- Deleting unallocated volumes on page 157
- Unallocating volumes from file servers on page 228
- Unallocating volumes from hosts on page 227

About deleting unallocated volumes

Volumes that are not allocated to any host can be deleted and their space added to the unused capacity of DP pools or parity groups. To completely and securely remove the data, shred the volume data before deleting the volume.

Related concepts

- About creating volumes on page 149
- About removing hosts and related resources, or vCenters on page 79
- About releasing a LUSE volume on page 159

Related tasks

- Deleting unallocated volumes on page 157

Deleting unallocated volumes

You can delete unallocated volumes from a registered storage system.

Before you begin

- Identify the target storage system
- Identify the target volumes
- Shred volume data, if needed
- Unallocate volumes

Procedure

1. On the **Resources** tab, select the target storage system.
2. Expand the tree and select the storage system from which you want to delete volumes.
3. Select **Open-Unallocated** or **DP Pools** and then select the **DP Vols** tab of the target DP pool.
4. From the volume list, select the volumes that you want to delete, and then click **Delete Volumes**.
5. Specify additional information, as appropriate:
   - Verify the information that is displayed.
   - Enter a name in **Task Name**.
   - Specify when to execute the task.
6. Click **Submit**, and confirm task completion.

Result

When the task completes, deleted volumes no longer appear in the Open-Unallocated or DP volume list.

Related concepts

- About deleting unallocated volumes on page 156

Related tasks

- Shredding volume data on page 155
- Unallocating volumes from file servers on page 228
- Unallocating volumes from hosts on page 227

About creating a LUSE volume

The functionality to create LUSE volumes is supported by certain mid-range and enterprise storage systems. However, mid-range storage systems do not ensure data integrity when creating LUSE volumes. Therefore, host administrators must back up data to a server or other location beforehand.

A LUSE volume is collection of two or more basic volumes that have been grouped together to increase available capacity.

If the capacity of the LUSE volumes or basic volumes that have already been allocated to hosts becomes insufficient, you can increase capacity by adding unallocated volumes or creating LUSE volumes.

LUSE volumes are created by using volumes that belong to the same resource group, have the same drive type, and are at the same RAID level. When using FMD/SSD flash drives in a Virtual Storage Platform or HUS VM storage system, you can still create a LUSE volume even if the drive types are different.

Related tasks

- Creating a LUSE volume on page 158

Creating a LUSE volume

LUN size expansion (LUSE) volumes are created by aggregating volumes into a larger logical volume.

Before you begin

- Gather the name of the target host
- Identify the target volumes (LUSE volumes or basic volumes that are already allocated to the host).
- Back up the target volume (for mid-range storage systems this includes HUS 100, AMS 2000, SMS, and AMS/WMS). Before creating LUSE volumes in mid-range storage systems, such as the AMS series, host administrators must back up data to a server or other location. This is because data might be corrupted when LUSE volumes are created.
- Create volumes to be added

Procedure

1. On the **Resources** tab, select **Hosts**.
2. Expand the tree and select the target OS.
3. Select the target host from the list of hosts, and from the displayed list of volumes, select the LUSE volume or basic volume, and then click **Create LUSE Volume**.
4. Click **Add Volumes** and select volumes to be added.
5. Click **Show Plan** and confirm that the information in the plan summary is correct. If changes are required, click **Back**.
6. (Optional) Update the task name and provide a description.
7. (Optional) Expand **Schedule** to specify the task schedule.

   You can schedule the task to run immediately or later. The default setting is **Now**. If the task is scheduled to run immediately, you can select **View task status** to monitor the task after it is submitted.
8. Click **Submit.**

If the task is scheduled to run immediately, the process begins.

9. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.

Result

Information about the created LUSE volume can be verified in the list of volumes for the target host on the Resources tab.

Related concepts

- About creating a LUSE volume on page 157

About releasing a LUSE volume

To release a LUSE setting, the LUSE volume must first be unallocated from all hosts. Users can release a LUSE setting either by selecting it or using the Unallocating Volumes dialog box.

Related concepts

- About removing hosts and related resources, or vCenters on page 79
- About deleting unallocated volumes on page 156

Related tasks

- Releasing a LUSE volume on page 159

Related references

- Unallocate volumes dialog box on page 229

Releasing a LUSE volume

Releasing the LUSE setting for a target volume releases the LUSE volume.

Before you begin

Before releasing a LUSE volume, make sure that the volume has been unallocated from all hosts.

Procedure

1. On the **Resources** tab, expand the **Storage Systems** tree, and for the target storage system, select **Volumes** and **Open-Unallocated**.
2. From the list of volumes, select the volume for which the LUSE setting is to be released, and click **Release LUSE Volumes**.
3. Verify that the information in the Plan Summary is correct.
4. (Optional) Update the task name and provide a description.
5. (Optional) Expand **Schedule** to specify the task schedule.

You can schedule the task to run immediately or later. The default setting is **Now**. If the task is scheduled to run immediately, you can select **View task status** to monitor the task after it is submitted.

6. Click **Submit.**
   If the task is scheduled to run immediately, the process begins.
7. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.

Result

When the task completes, the LUSE setting for the target volume is released.

Related concepts

- About releasing a LUSE volume on page 159

Virtualizing external storage

This module describes virtualizing external storage so that external volumes can be managed.

About virtualizing and unvirtualizing volumes

Virtualizing volumes is the process of mapping external storage system volumes to an internal storage system.

Virtualized volumes allow for centralized management of volumes from multiple external storage systems, in addition to the volumes from the single storage system being managed with Hitachi Command Suite.

The volumes in the external storage system (registered and referred to as *external volumes*) can be managed in the same way as the volumes in the internal storage system (registered and referred to as *internal volumes*). Volume virtualization provides the capability to optimize storage services to applications by facilitating movement of data between tiers and migration of data between different generations of technology.

Note: Internal and external storage systems can be connected by Fibre Channel, Fibre Channel over Ethernet, or iSCSI. Connections between different port types are not supported.

When external storage is managed by an internal storage system, external volumes will be listed in the External Storage folder of the internal storage system.

External volumes can be used for various purposes:

- External volumes can be allocated to hosts the same as internal volumes are.
- Important data on internal volumes can be backed up to external volumes.

- Infrequently accessed data on internal volumes can be moved to external volumes for archiving.

You can virtualize volumes in two ways:

- Virtualizing volumes:
  In HCS, external volumes are mapped by specifying the external storage system, internal storage system, external volumes, and number of external paths. Although HCS manages only registered storage systems, external volumes in unregistered storage systems can also be mapped.
  If the external storage system is registered in HCS, you can directly specify individual external volumes in that storage system, or you can specify just the size and number of volumes. HCS suggests mapping plans accordingly.
- Discovering and virtualizing volumes:
  If HCS does not manage the external storage system, HCS detects the external volumes allocated to the internal storage system's external ports, and then selects the volume to be used. If no volume that meets the conditions exists in the external storage system, a volume is automatically created from free space in the parity group.

When you want to perform a data migration or discard devices, you must cancel the virtualization. When you unvirtualize volumes, the virtualized volumes are unmapped from the specified external storage.

- If you are unvirtualizing volumes from a registered storage system, the volumes are automatically unallocated. You also have the options to delete the volumes or host group.
- If you are unvirtualizing volumes from an unregistered storage system, the volumes are only unmapped. Other management tasks are not available.

Tip: For VSP Fx00 models, VSP Gx00 models or VSP G1000, VSP G1500, or VSP F1500 storage systems, you can reconnect or disconnect external storage systems, reconnect or disconnect external volumes, edit path configurations, or unvirtualize external volumes by using the windows available by clicking the System GUI link. To access the System GUI link, on the Resources tab, right-click External Storage for the target storage system, and then select System GUI from the menu. Or, click External Storage for the target storage system, and then click the System GUI link that appears in the application pane.

For information about how to perform tasks that are available by clicking System GUI, see the Help for the appropriate window or see the *Provisioning Guide* for your storage system.

For related information about external storage virtualization, see the *Hitachi Universal Volume Manager User Guide*.

Related concepts

- Creating a storage operating environment on page 140

Related tasks

- Virtualizing volumes of a registered storage system on page 162
- Discovering and virtualizing volumes of an unregistered storage system on page 163
- Unvirtualizing volumes on page 165

Virtualizing volumes of a registered storage system

Choose from a list of existing unallocated volumes in the external storage system to be virtualized. Or specify the number of volumes and volume capacity to search for volumes that meet the criteria.

Before you begin

- The internal and external storage systems are registered in Hitachi Command Suite.
- The external storage system is connected to the internal storage system by Fibre Channel or Fibre Channel over Ethernet (FCoE).
- Hitachi Universal Volume Manager is installed on the internal storage system.
- The port attribute settings for internal storage system.
- Specify the volumes or identify the number and capacity of volumes you are virtualizing.
- Identify the number of external paths between the external and internal storage systems. The maximum number of external paths that can be set is eight.

Procedure

1. From the **Actions** menu or **General Tasks** pane, select **Virtualize Volumes**.
2. From the list, select an external storage system. A storage system is preselected if you selected a storage system from the Storage Systems tree and launched the dialog box.
3. From the list, select an internal storage system.
4. Choose the volumes by clicking one of the following options:
   - **Select from existing available volumes**: Click **Add/Remove Volumes** to select from a list of available volumes, then click **Add** to add them to the Selected Volumes table.
   - **Specify number of volumes and volume capacity**: From the list, select the **No. of Volumes** and enter the **Volume Capacity**. Click **Advanced Options** to specify drive type, drive speed, RAID level, or parity group in the volume criteria. The storage system will search for volumes that match the criteria.
5. Specify the **No. of External Paths per Volume**.
6. (Optional) Click **External Path Priority Settings** to configure the external path and increase or decrease external path priority.

7. (Optional) Click **Host Group and LUN Settings** to configure the host group settings and LUN number.
8. (Optional) Click **Internal Volume Configuration** to configure the options for the volume to be virtualized, such as the parity group ID, CLPR, inflow control, LDEV labels, and LDEV IDs.
9. Click **Show Plan** and confirm that the information in the plan summary is correct. If changes are required, click **Back**.
10. (Optional) Update the task name and provide a description.
11. (Optional) Expand **Schedule** to specify the task schedule.

    You can schedule the task to run immediately or later. The default setting is **Now**. If the task is scheduled to run immediately, you can select **View task status** to monitor the task after it is submitted.
12. Click **Submit.**
    If the task is scheduled to run immediately, the process begins.
13. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.

Result

When the task completes, the existing volumes are allocated to the internal storage system and virtualized.

Tip: When the task to virtualize volumes fails, if external paths are already set but have not been mapped to volumes in the internal storage system, from the external storage systems drop-down list in the Virtualize Volumes dialog box, select Search new systems/volumes by port scan to map the paths.

Related concepts

- About virtualizing and unvirtualizing volumes on page 160

Related tasks

- Unvirtualizing volumes on page 165
- Installing the software on page 49

Discovering and virtualizing volumes of an unregistered storage system

Use the port scan and volume discovery functionality to detect external storage system volumes allocated to the external port and then virtualize the available volumes.

Note: To virtualize a volume that uses an iSCSI port, in the Resource tab, right-click External Storage under the target storage system, and then select System GUI. Alternatively, click the System GUI link that displays in the applications area. Then, virtualize the volume. For more information about the System GUI, see the *Hitachi Universal Volume Manager User Guide*.

Before you begin

- The internal storage system is discovered in Hitachi Command Suite.
- The internal storage system is VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, VSP F1500, Virtual Storage Platform, or HUS VM that has the port scan and volume discovery functionality.
- The external storage system is connected to the internal storage system by Fibre Channel or Fibre Channel over Ethernet (FCoE).
- Universal Volume Manager is installed on the internal storage system.
- The port attribute settings for internal storage system.
- The external mapping volumes are already allocated to the external ports.

Procedure

1. From the **Actions** menu or **General Tasks** pane, select **Virtualize Volumes**.
2. From the list of external storage systems, select **Search new systems/ volumes by port scan**.
3. From the list, select an internal storage system.
4. Specify the external paths by choosing one of the following options:
   - Click **Select newly discovered external paths** and continue with step 5.
   - Click **Select an existing external path group**, select an external path group from the list, and continue with step 7.
5. Click **Start Port Scan**.
   The external storage is discovered and the port ID and external WWN are displayed.

   Note: Because the external storage is discovered, you can virtualize additional volumes at a later time without doing a port scan, provided that the same external paths are used. If you are virtualizing additional volumes at a later time and not using the same external paths, then you need to use the port scan functionality to discover available external storage volumes.

6. Select an external storage, then choose the external paths to be used for virtualization and click **OK**.
7. (Optional) Click **External Path Priority Settings** to configure the external path and increase or decrease external path priority.
8. Click **Start Volume Discovery**.
   The volumes that can be accessed using the selected external paths are displayed.
9. Review the list of volumes to be virtualized and click **OK**.
10. (Optional) Click **Internal Volume Configuration** to configure the options for the volume to be virtualized, such as the parity group ID, CLPR, inflow control, LDEV labels, and LDEV IDs.

11. Click **Show Plan** and confirm that the information in the plan summary is correct. If changes are required, click **Back**.
12. (Optional) Update the task name and provide a description.
13. (Optional) Expand **Schedule** to specify the task schedule.
    You can schedule the task to run immediately or later. The default setting is **Now**. If the task is scheduled to run immediately, you can select **View task status** to monitor the task after it is submitted.
14. Click **Submit.**
    If the task is scheduled to run immediately, the process begins.
15. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.

Result

When the task completes, the volumes on the discovered external storage system are virtualized and can be managed from HCS.

Related concepts

- About virtualizing and unvirtualizing volumes on page 160

Related tasks

- Unvirtualizing volumes on page 165
- Installing the software on page 49

Unvirtualizing volumes

When you unvirtualize volumes, virtualized volumes are unmapped from the specified external storage system.

Note: To unvirtualize a volume that uses an iSCSI port, in the Resource tab, right-click External Storage under the target storage system, and then select System GUI. Alternatively, click the System GUI link that displays in the applications area. Then, unvirtualize the volume. For more information about the System GUI, see the *Hitachi Universal Volume Manager User Guide*.

Before you begin

- Volumes to be unmapped on the storage system are virtualized.

Procedure

1. From the **Actions** menu or **General Tasks**, select **Unvirtualize Volumes**.
2. From the list, select an external storage system.
3. From the list, select an internal storage system.
   The parity groups that correspond to the virtualized volumes from the selected external storage system are displayed in the Selected Parity Groups table. Only parity groups that are available to unmap are listed in the table.

4. (Optional) Click **Add/Remove Parity Groups**. In the Selected Parity Groups table, select any external parity groups that you do not want unmapped and click **Remove** to move them to the Existing Parity Groups table. You can select several external parity groups at one time.
5. (Optional) Click **Advanced Options** to configure deleting the host group, releasing LUSE volumes, or deleting volumes.

   Note: If the selected external storage system is not a registered storage system, you can unmap volumes only.

6. Click **Show Plan** and confirm that the information in the plan summary is correct. If changes are required, click **Back**.
7. (Optional) Update the task name and provide a description.
8. (Optional) Expand **Schedule** to specify the task schedule.

   You can schedule the task to run immediately or later. The default setting is **Now**. If the task is scheduled to run immediately, you can select **View task status** to monitor the task after it is submitted.
9. Click **Submit.**
   If the task is scheduled to run immediately, the process begins.
10. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.

Result

When the task completes, the volumes on the selected external storage system are unmapped.

Related concepts

- About virtualizing and unvirtualizing volumes on page 160

Related tasks

- Virtualizing volumes of a registered storage system on page 162
- Discovering and virtualizing volumes of an unregistered storage system on page 163

Virtualize Volumes dialog box

Volume virtualization provides the capability to optimize storage services to applications by facilitating movement of data between tiers and migration of data between different generations of technology.

When you enter the minimum required information in a dialog box, the Show Plan button activates to allow you to review the plan. Click the Back button to modify the plan to meet your requirements.

The following table describes the dialog box fields, subfields, and field groups. A *field group* is a collection of fields that are related to a specific

action or configuration. You can minimize and expand field groups by clicking the double arrow symbol (>>).

As you enter information in a dialog box, if the information is incorrect, errors that include a description of the problem appear at the top of the box.

Table 11 Virtualize Volumes dialog box

Field	Subfield	Description
Select the external and internal storage systems	-	External storage system is displayed if you selected a registered storage system and launched the dialog box. If you did not, select an external storage system from the list. From the list of external storage systems, select Search new systems/volumes by port scan to scan ports and discover volumes of unregistered storage systems. From the list of internal storage systems, select the internal storage system that you want to use.
	Start Port Scan	Available when Search new systems/volumes by port scan is selected. Before clicking Start Port Scan, you have the options for specifying the external paths by choosing Select newly discovered external paths or Select an existing external path group.
	Add/Remove External Paths	After the port scan, the discovered external storage is displayed in alphabetical order. The first one is selected by default. The newly discovered external paths between the selected internal and external storage systems are listed. Click Add to add the external paths that you want to virtualize.
	Start Volume Discovery	Click to search for available volumes.
Select from existing available volumes	-	Select and click Add/Remove Volumes. From the list of available external volumes, select the volume to be allocated to the internal storage system and click Add.
Specify the number of volumes and volume capacity	-	Specify the Number of Volumes that you are virtualizing. The new volumes will be partitioned on the external storage system and then allocated to the internal storage system. In the Volume Capacity dialog box, enter the storage capacity of the volumes you are virtualizing. Specify the value in MB, GB, or TB. The required capacity is based on the number of volumes and the volume capacity.
>> Advanced Options	(See the following four fields for details)	Click >> Advanced Options to display and configure drive and parity group options.
Drive type	-	Existing drive type is displayed based on the selected storage system. If multiple drive types are displayed, accept the default or select from the drive type list:

Field	Subfield	Description
		• Any • FC - Fibre channel • SAS - Serial Attached SCSI • SATA - Serial ATA • FMC - Flash Memory Compressed • FMD - Flash module drive • SSD - Solid state drive • SAS (SED) - SAS self-encrypting drive Any is displayed when you can select more than two drive types. Changing drive type can change the parity group.
Drive speed (RPM)	-	Existing drive speed is displayed based on the selected storage system. If multiple drive speeds are displayed, accept the default of any available speed or select from the drive speed list: • Any • 15000 • 10000 • 7200 Any is displayed when you can select more than two drive speeds. Changing the drive speed (rotational RPM) can change the parity group.
RAID level	-	Existing RAID levels are displayed for the external storage system. For example, RAID 0, 1+0, 5, and 6 are valid RAID configurations when 6 drives are available (for example, HUS 100). If the HUS 100 storage system has RAID 5 (3D+1P) parity groups, then RAID 5 (3D+1P) is displayed. The RAID levels change in the dialog box based on the selected storage system and supported RAID levels.
Parity Group	Select Parity Group	The parity group. Based on the drive type, drive speed, and RAID level values, an appropriate parity group is selected. Select a parity group by clicking Select Parity Group. From the list of parity groups, sort and filter columns such as RAID level, unallocated capacity, or other fields to identify the preferred parity groups.
No. of External Paths per Volume	-	Select the number of external paths between the external and internal storage systems. The number can be 1 through 8.
External Paths	Ext. Storage Port	The port of the connected external storage system. This port is a part of the external storage and used to connect to the internal storage.
	Controller	The external storage controller device information. This device controls the transfer of data between an external storage system and an internal storage system.
	Ext. Storage System	The connected external storage system.

Field	Subfield	Description
	Int. Storage System	The connected internal storage system.
	Int. Storage Port	The port of the connected internal storage system. This port is a part of the internal storage and used to connect to the external storage.
	Edit External Paths	Edit external paths from a table or from a graph. To create an external path, select an external storage port and an internal storage port. To delete an external path, select the path that you want to delete. Click Clear External Paths. External paths can be edited from a table or from a graph. To create an external path, select an external storage port and an internal storage port. To delete an external path, select the path that you want to delete. To delete all paths, click Clear External Paths.
	Port Type (Ext./Int.)	The port types of the storage ports of both the external and internal storage systems are displayed.
>> External Path Priority Settings	(See the following four fields for details)	Click >> External Path Priority Settings to increase or decrease the external path priority.
Ext. Storage Port	-	The port for the connected external storage system.
Int. Storage Port	-	The port for the connected internal storage system.
Port Type (Ext./ Int.)	-	The storage port type of the external path. The port types of the storage ports of both the external and internal storage systems are displayed.
Increase Priority	-	Increase the priority of the selected external path. The path with the highest priority is the primary path if the external storage system works in a single path mode.
Decrease Priority	-	Decrease the priority of the selected external path. The path with decreased priority is the alternate path if the external storage system works in a single path mode.
>> Host Group and LUN Settings	(See the following four fields for details)	Volume allocations using Fibre Channel and FCoE prompt for such items as host group name, host mode, host mode options, and LUN number. To display or hide the following fields, click Host Group and LUN Settings.
Host group name	-	Specify a host group name for the internal storage system. Paths are allocated from the external storage system to the internal storage system.
Host Group Settings in Detail	-	The host groups to be used during volume allocation or that were previously set during prior volume allocation.
Host Mode Options	-	Select one or more host mode options for supporting special requirements for specific applications.
LU Number (hex)	-	Assign LU numbers to the allocated volumes by entering a Start from number (or using the default). Set LU numbers in ascending order and do not use existing numbers.
>> Internal Volume Configuration	(See the following four fields for details)	Click >> Internal Volume Configuration to display and configure the internal volume options.

Field	Subfield	Description
Parity Group	-	Assign internal volumes to a parity group by entering a Start from number. For Device Manager, when you specify the first number (1-16384), the second number is automatically assigned. The format for specifying external parity groups is, for example, "E1" or "E16384".
CLPR	-	From the list, select the cache logical partition resource (CLPR) value for the internal volumes. If you are using a VSP G400, VSP G600, or VSP G800 storage system with a NAS module and the firmware version is earlier than 83-03-25-XX/XX, do not use the CLPR named "NASSystemCLPR". If you use this CLPR, system performance might be severely affected.
Inflow control	-	Select Enable to stop write operations to cache and to stop I/O from the host when the write operations to the external volume fails. Select Disable to write I/O from the host during retry operations to cache memory, even when the writing operation to the external volume fail. When write operations to the external volume returns to normal, all the data in cache memory is written to the external volume. The default is Disable.
Cache mode	-	Select Enable, after receiving the data into the internal storage system cache memory, the internal storage system signals the host that an I/O operation has completed and then asynchronously writes the data in cache memory to the external storage system. If you select Disable, the internal storage system signals the host that an I/O operation has completed only after the local storage system has synchronously written the data to the external storage system. The default is Enable.
Resource group	-	Resource groups listed by storage system, parity group, LDEV ID, and storage port. Resource group appears only when multiple resource groups are allocated and each resource group has LDEV IDs of internal volumes that belong to the selected internal storage system.
Label	-	Volume labels are searchable, and therefore recommended as a way to find volumes. Specify and add a LDEV label.
Initial value	-	The smallest number of sequential numbers. The Initial value is not required but can be useful for differentiation when virtualizing multiple volumes. For each volume, the number is in ascending order of LDEV ID.

Field	Subfield	Description
Reflect these labels to the storage system	-	Reflect these labels to the storage system is checked by default so that naming is consistent between HCS and the storage system. If the selected storage system does not support label setting, this item is not displayed.
Quorum Disk	-	When virtualizing a Quorum disk for use in a global-active device environment, Change Internal Volume(s) to Quorum Disk(s) must be selected. The Remote Storage System is displayed, and the recommended disk ID displayed, but can be changed if necessary. Use Read Response Guaranteed Time When Quorum Disk Blocked to set the time required to change the status of global-active device pair volumes to `PSUE (Block)` when failures occur in the quorum disk and also the primary or secondary storage system. To change this value after creating a quorum disk, use Device Manager - Storage Navigator.
LDEV ID	-	LDEV ID options include auto or manual.

Virtualizing storage capacity (HDP/HDT)

This module describes how to virtualize storage capacity.

About virtualizing storage capacity

Dynamic Provisioning (HDP) provides virtual volumes to a host and allocates the actual capacity from a HDP pool when a host makes a write request. By using HDP pools, you can allocate more capacity to a host than that allowed by the actual physical configuration of the storage system.

HDP pools provide the following advantages:

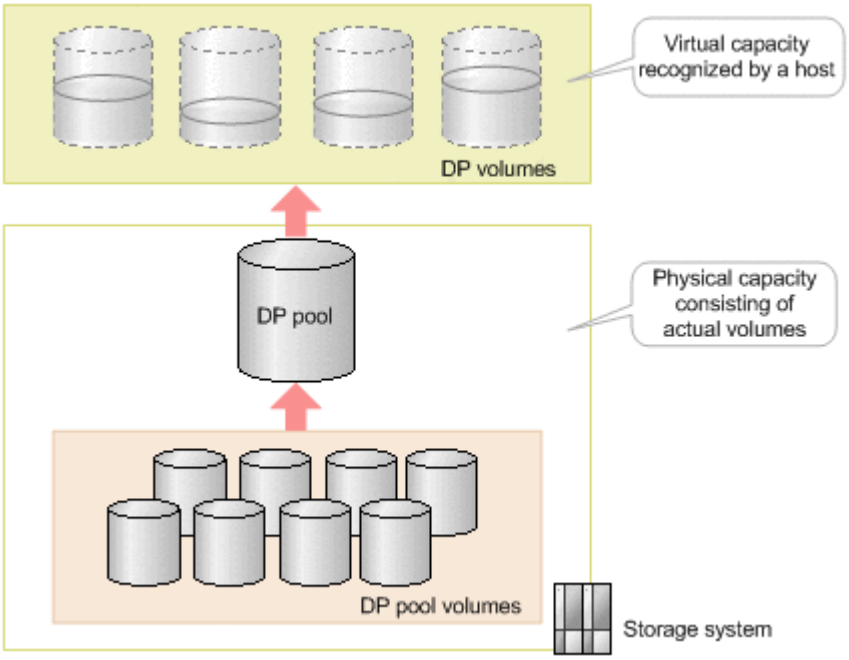
- You can reduce system setup and operational costs.
- You can use resources more efficiently.
- You can distribute workload equally among volumes.

Dynamic Tiering (HDT) improves HDP pools storage performance further by taking advantage of performance differences between hardware tiers.

Note: The terms "HDP" and "HDT" are referred to collectively as "DP".

You can create virtual volumes (DP volumes) from physical volumes that are grouped into DP pools. You can then allocate those virtual volumes to hosts.

In the illustration below, note that the volumes used to create the DP pool are called DP pool volumes. The DP pool is then used to provide capacity as needed to allocated DP volumes.



DP pools are created by Device Manager, which automatically selects volumes based on user-specified conditions. You can also directly specify parity groups to create a DP pool.

By enabling the Full Allocation settings, you can reserve pages of the specified capacity in advance (referred to as "reserved capacity"). The reserved capacity is included in the used capacity. You can also specify values for the usage rate threshold and reservation threshold.

If you have registered a Tiered Storage Manager license, you can use the Mobility tab to evaluate and analyze the operation status that is related to a DP pool. When you delete an unnecessary DP pool, DP volumes created from that pool are also deleted. For this reason, a prerequisite to deleting a DP pool is that no DP volumes from that pool are allocated to a host.

Tip: For VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, or VSP F1500 storage systems, you can restore DP pools, complete SIMs, and export tier relocation log files by using the windows available by clicking the System GUI link. To access the System GUI link, on the Resources tab, right-click DP Pools for the target storage system, and then select System GUI from the menu. Or, click DP Pools for the target storage system, and then click the System GUI link that appears in the application pane.

For information about how to perform tasks that are available by clicking System GUI, see the Help for the appropriate window or see the *Provisioning Guide* for your storage system.

For Unified Storage 100 family of storage systems, the replication data and replication management area, which are used by Copy-on-Write Snapshot or TrueCopy Extended Distance, are stored in the created DP pool.

Related concepts

- Creating a storage operating environment on page 140

Related tasks

- Creating a DP pool on page 173
- Verifying DP pool information on page 180
- Expanding DP pools on page 180
- Shrinking a DP pool on page 183
- Modifying DP pool settings on page 184
- Deleting DP pools on page 184
- Expanding DP volumes on page 185
- Reclaiming zero pages on page 186

Related references

- Create Pool dialog box on page 176
- Conditions for DP pools that are used in storage pools on page 253

Creating a DP pool

You can create an HDP or HDT pool, which provides more efficient use of physical storage for virtual volumes that are allocated to hosts from the DP pool. DP pool performance can be improved if you use the entire capacity of a parity group for a single DP pool.

Before you begin

- Register the target storage system.
- When defining an external LDEV tier rank, externally connect a storage system that has multiple performance levels.
- Verify the following when using VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, VSP F1500, Virtual Storage Platform, Hitachi Universal Storage Platform V/VM , or Hitachi Unified Storage VM (HUS VM) storage systems for DP pools:
  - Parity groups from which volumes have already been created can be added
  - Drive type characteristics (for example: drive type and drive speed) in parity groups and RAID level
- Verify the following when using HUS 100 and the Adaptable Modular Storage (AMS) 2000 family of storage systems for DP pools:
  - Drive type and RPM (only drives in which a parity group has not been created can be targets)
  - Parity group RAID level and capacity
  - Number of parity groups

Note: In HDT pools, if different drive types and/or RAID levels are mixed in a single tier, they will all be considered equal for data placement regardless of page access frequency. As a result, I/O performance will be dependent on the drive type characteristics and RAID level on which any given page resides.

In HDP pools, if different drive types and/or RAID levels are mixed in an HDP pool, I/O performance will be dependent on the drive type characteristics and RAID level on which any given page resides.

Procedure

1. On the **Resources** tab, expand the storage system, list existing **DP Pools**, and click **Create Pool**.
2. In the **Create Pool** dialog box, specify a pool name, and optionally select **Reflect this pool name to the storage system**.
3. To configure an HDP pool, select a **Pool Type** of **HDP** and configure the following:
   a. In the **Additional Parity Groups** table, click **Add Parity Groups**.
   b. (Optional) Select 'Allow to mix different drive types/speeds, chip types, RAID levels or volume locations' to allow combining resources with different characteristics.
   c. Select one or more parity groups, and click **Add to Pool** and then **Close**. The **Pool Summary** information is updated.
   d. (Optional) Click **Advanced Options** to configure Pool ID, Used Threshold, Subscription Thresholds, and DP volume protection options, as needed.
   e. Go to step 7.
4. To configure an HDT pool, select a **Pool Type** of **HDT**, and then in the **Additional Parity Groups** table, choose **Standard** or **Mixed** mode to disallow or allow combining resources with different characteristics. For **Mixed** mode, go to step 6.
5. For a **Standard** mode HDT pool, do the following:
   a. Click **+** to add a new tier.
   b. In the **Add New Tier** dialog box, select a volume to configure Tier 1, and click **Select**. The **Tier Configuration** table in **Pool Summary** is updated.
   c. Click **Add Parity Groups**, select the parity group, click **Add to Pool**, and click **Close**. Select an available parity group that best meets your performance or capacity needs (Tier 1 for best performance, Tier 2 for next best performance, and Tier 3 for capacity).
   d. (Optional) Click **+** to add Tier 2 and Tier 3, configure the tiers based on your performance and capacity needs using the choices in **Add New Tier** dialog box. The Tier Configuration table in Pool Summary is updated.

   Tip: To delete an existing tier, click **X** in the **Tier** tab.

e. (Optional) Click **Advanced Options** to configure Pool ID, Used Threshold, Subscription Thresholds, and DP volume protection options, as needed.
f. Click **HDT Options**, and configure the tier management options as needed.
g. Go to step 7.

6. For a **Mixed** mode HDT pool, do the following:
   a. Click **Add Parity Groups**.

   Note: There are two parity group tab choices from which you can select: **Internal Parity Groups** and **External Parity Groups** (the **Internal Parity Groups** tab is set by default). If you select the **External Parity Groups** tab and select one or more parity groups, this enables the **External LDEV Tier Rank** menu from which you must choose a ranking for the tier.

   b. For mixed mode in the **Internal Parity Groups** tab or the **External Parity Groups** tab, select parity groups that you want to add to the HDT pool, click **Add to Pool**, and click **Close**. The Tier Configuration table in Pool Summary shows the new tier configuration status for each tier.
   c. (Optional) Click **Advanced Options** to configure Pool ID, Used Threshold, Subscription Thresholds, and DP volume protection options, as needed.
   d. Click **HDT Options**, and configure the tier management options as needed.
7. Click **Show Plan** and confirm that the information in the plan summary is correct. If changes are required, click **Back**.
8. (Optional) Update the task name and provide a description.
9. (Optional) Expand **Schedule** to specify the task schedule. You can schedule the task to run immediately or later. The default setting is **Now**.
10. Click **Submit**. If the task is to run immediately, the task begins.
11. You can check the progress and the result of the task on the **Tasks & Alerts** tab. Click on the task name to view details of the task.

Result

Created pools are added to the target storage system DP Pools list.

Related concepts

- About virtualizing storage capacity on page 171

Related tasks

- Verifying DP pool information on page 180
- Modifying DP pool settings on page 184
- Deleting DP pools on page 184

Related references

- Create Pool dialog box on page 176

Create Pool dialog box

Pools can be created for storage systems that support HDP. In addition, HDT pools of differing performance levels can be used to improve application performance.

When you enter the minimum required information in this dialog box, the Show Plan button activates to allow you to review the plan. Click the Back button to modify the plan to meet your requirements.

The following table describes the dialog box fields, subfields, and field groups. A *field group* is a collection of fields that are related to a specific action or configuration. You can minimize and expand field groups by clicking the double-arrow symbol (>>).

As you enter information in a dialog box, if the information is incorrect, errors that include a description of the problem appear at the top of the box.

Table 12 Create Pool dialog box

Field	Subfield	Description
Storage System	-	Displays the selected storage system name, or prompts you to select the storage system from a list.
Pool Name	-	Accept the default pool name, or enter a pool name. Do not confuse pool name with pool ID. The pool ID is an assigned number, the pool name is a user-definable value. Reflect a pool name to the storage system is selected by default, and provides naming consistency between HCS and the storage system. If it is not displayed, it does not apply to the selected storage system.
Pool Type	HDP or HDT	If Pool Type with HDP and HDT options is displayed, the selected storage system supports both HDP and HDT. If Pool Type options are not displayed, it means the selected storage system only supports HDP. Select HDP to create a pool using one or more parity groups. Select HDT to create a pool using one or more tiers, each with one or more parity groups. The HDT pool type offers two mode option choices, Standard or Mixed, and HDT options can be minimized or expanded at the bottom of the dialog box.
Pool Summary	-	Pool Summary information is updated as parity groups are selected (number of parity groups, DP pool capacity, used capacity % (for thresholds 1 and 2), and subscription % (that trigger subscription warnings and limits). When pools are configured, use Advanced Options to set the Used Thresholds 1 and 2, and the Subscription thresholds for warnings and limits. For HDP, Pool Configuration displays the physical attributes of the pool, including volume location, drive type, drive speed, chip type, RAID level, number of parity groups, and parity group capacity.

Field	Subfield	Description
		For HDT, Tier Configuration displays the attributes of one or more tiers, including tier number (1-3), volume location (internal/external), drive type, drive speed, chip type, external LDEV tier rank, number of parity groups, and parity group capacity.
Additional Parity Groups	-	(Information in the Additional Parity Groups table differs depending on the selected pool type.) For an HDP pool, this table displays parity group, drive type, drive speed, chip type, RAID level, total capacity, unallocated capacity, free capacity, number of available volumes, external storage system, external storage model, and cache mode. For an HDP pool, you can select Allow mixing of different drive types/speeds, chip types, RAID levels, or volume locations to create pools using parity groups with different attributes. Click Add Parity Groups to add one or more parity groups to a pool, and click Close to review the selected parity groups. Click Remove Parity Groups to remove any parity groups from the list. For an HDP pool, the Available Parity Groups table lists parity group, drive type, drive speed, chip type, RAID level, total capacity, unallocated capacity, free capacity, number of volumes, CLPR, cache mode, and resource group. For an HDT pool, when the Internal Parity Groups tab is selected, the Available Parity Groups table lists parity group, drive type, drive speed, chip type, RAID level, total capacity, unallocated capacity, free capacity, number of volumes, CLPR, cache mode, and resource group. When the External Parity Groups tab is selected, the Available Parity Groups table lists parity group, total capacity, unallocated capacity, free capacity, number of volumes, CLPR, external storage system, external storage model, cache mode, and resource group. If the External Parity Groups tab is selected and you select one or more parity groups, the External LDEV Tier Rank menu is enabled from which you must choose a ranking for the tier. For an HDT pool, select from two mode option choices; Standard (single drive type/speed, chip type, RAID level, and volume location in a single tier) or Mixed (mixed drive types/speeds, chip types, RAID levels, and/or volume locations in a single tier). The selected mode affects the available parity groups you can add to pools in each tier to best meet your performance or capacity needs (Tier 1 for best performance, Tier 2 for next best performance, and Tier 3 for capacity). (For Standard mode) Click + to display the Add New Tier dialog box that displays volume location, drive type, drive speed, and external LDEV tier ranking. Select a volume to add to the tier (for example, Internal, SAS, 15000 RPM), and click Select to update Tier Configuration with the selected volume and drive information, and create a Tier 1 tab. On the Tier 1 tab, use Add Parity Groups to add parity groups to this tier. Click + to create Tier 2, and again to create Tier 3. (For Mixed mode) Click Add Parity Groups, select parity groups from Internal Parity Groups or External Parity Groups tab (for example: select the check box for the title row in Internal Parity Groups, click Add to Pool, and click Close). The Tier Configuration table in Pool Summary displays the new status for each configured tier. Note the following important items regarding HDT configuration: • HCS automatically arranges your tiers from highest performance (Tier 1) to lowest performance (Tier 3), regardless of the order used when creating

Field	Subfield	Description
		tiers. For example, if Internal, SAS, 15000 is the tier created, but it is the highest performance tier, then Internal, SAS, 15000 is displayed as Tier 1. • When three tiers are defined, the + tab is no longer displayed because there is a three-tier limit. To delete a tier, click X in the tier tab. When the + tab displays, you can define a new tier. Create new tiers and delete existing tiers as required.
>> Advanced Options	Pool ID	Accept the default pool ID number, or choose from the options in the menu. Use Show Pool ID Usage to display pool ID number and names that are currently in use.
	Used Threshold 1	Set the level of physical capacity usage capacity as a percentage (1-100 range) of the physical DP pool capacity. When this threshold is exceeded, it generates an alert, an email, or both. Alerts display on the Dashboard and the Tasks & Alerts tab.
	Used Threshold 2	Same description as above. Used Threshold 1 and 2 display in the Pool Summary, in both graphic and numeric values.
	Subscription Thresholds	(Optional) Select Enabled to set subscription warning and limit thresholds.
	Warning	Exceeding the subscribed warning threshold (as a percentage of DP pool capacity) generates an alert, an email, or both. Warnings are generated when volume allocation from a pool will exceed the warning level, but volume allocation is still allowed (this is a soft limit). Alerts display on the Dashboard and the Tasks & Alerts tab.
	Limit	Exceeding the subscribed limit threshold (as a percentage of DP pool capacity) generates an alert, an email, or both. However, if volume allocation will exceed the subscribed limit threshold, volume allocation is not allowed (this is a hard limit). You must either reduce the capacity of the volume you want created, or increase the subscription limit.
	Protect DP VOL when:	For the VSP G1000, VSP G1500, and VSP F1500 storage system, this option is displayed when: • The Data Retention Utility is licensed • The storage system microcode version is 80-02-01-XX/XX or later
	I/O fails to a full Pool:	Select 'Yes' to prohibit host I/O to one or more DP volumes in a DP pool that is full, where additional free space cannot be assigned. The displayed default (yes or no) is determined by values that have been set for the storage system. See DP pools to verify the protection settings of the DP pool. For DP volumes, see the volume information summary in volume details, and select the Guard link to verify that the protect attribute is assigned. If host I/O to DP volumes is suspended, resolve the issue, and then change the volume access attribute to read/write (see related topics).
	I/O fails to a blocked Pool VOL:	Select 'Yes' to prohibit host I/O to one or more DP volumes when a DP pool volume is blocked. The displayed default (yes or no) is determined by values that have been set for the storage system. See DP pool details to verify the protection settings of the DP pool. For DP volumes, see the volume information summary in volume details, and select the Guard link to verify that the protect attribute is assigned. If host I/O to DP volumes is suspended, resolve the issue, and then change the volume access attribute to read/write (see related topics).

Field	Subfield	Description
>> HDT Options	Tier management Auto	(HDT options are only displayed when it is this pool type.) By selecting this option, the storage system automatically starts and stops performance monitoring and tier relocation based on the Cycle time and Monitoring period settings. Cycle time values range from 30 minutes to 24 hours. Monitoring period are user-configurable time periods (for example, 13:00 to 19:00 hours). Continuous and Periodic monitoring modes apply to both the Cycle time and Monitoring period settings (for details, see Continuous Monitoring mode).
	Tier management Manual	Selecting this option lets you manually start and stop performance monitoring and tier relocation.
	Tier management Custom	Selecting this option lets you create performance monitoring and tier relocation plans using templates with different types of performance monitoring and tier relocation settings. You can create and delete templates by name, set up monitoring periods (up to 7 days a week), set up relocation start times (up to 7 days a week), and view these plans in the Schedule Summary table. You can modify templates to meet specific needs. In addition, click Pools using this template to assign a template to any pool that you select from the Pools table.
	Monitoring mode Continuous	Continuous monitoring mode uses weighted-average performance data gathered over several monitoring cycles so that tier relocation does not respond immediately to workload changes (see Cycle time). This mode prevents overreaction to workload changes, that might otherwise result in unnecessary tier relocation I/O.
	Monitoring mode Periodic	Periodic monitoring mode uses performance data from the last monitoring cycle (see Cycle time). Tier relocation responds immediately to workload changes. This mode is a more aggressive tier relocation monitoring mode.
	Relocation Speed	For VSP Gx00 models, VSP G1000, VSP G1500, or VSP F1500, you can select one of five speeds for the relocation of pages in a pool, in a unit of time, from 1 (slowest) to 5 (fastest). The default is 3 (standard). To reduce the load on the storage system, specify a slower page relocation speed. To give a higher priority to page relocations for the pool, specify a faster relocation speed.
	Buffer Space for New page assignment	Sets the buffer space used for new page assignments to HDT tiers (using the default values is recommended).
	Buffer Space for Tier relocation	Sets the buffer space used for tier page relocations between HDT tiers (using the default values is recommended).

Note: After configuring the pools, click Show Plan and perform all of the confirming, scheduling, and submitting tasks needed to create the pools.

Related concepts

- About virtualizing storage capacity on page 171

Related tasks

- Creating a DP pool on page 173
- Changing an access attribute to read/write on page 186

Verifying DP pool information

You can check the DP pools summary information to verify the current number of pools, the virtual volume capacity, the pools usage, and the total and used pool capacity.

Before you begin

- Create pools
- Discover (register) pools

Procedure

1. On the **Resources** tab, select **DP Pools** under the target storage system.

   A list of DP pools, (DP pool volumes and DP volumes) provides summary information that includes the number of pools, the virtual volume capacity, pool usage, and the total and used pool capacity.
2. Select a DP pool to display more detailed information about that pool.

   Use the detailed pools information to verify expected changes in capacity or performance. Configuration information for each DP volume is displayed by clicking the volume link from the list in the **DP Vols** tab.

   If you configured HDT pool tiers to be managed automatically using Element Manager, you can also verify the following:
   - Monitoring Information
   - Last Monitoring Start Date
   - Last Monitoring End Date
   - Monitoring Status
   - Relocation Status
   - Relocation Progress (%)

Related concepts

- About virtualizing storage capacity on page 171

Related tasks

- Creating a DP pool on page 173

Expanding DP pools

Pool capacity can be increased by expanding DP pools or by converting HDP pools into HDT pools.

Expand a DP pool by adding volumes to it. Expand an existing HDP pool by converting it into an HDT pool, which changes the pool type and allows you to configure tiers that support a mix of drive types and RAID levels.

Before you begin

- Register the storage system.

- Verify the external connections for any storage system with multiple performance levels before defining an external LDEV tier rank.
- Verify the following when using VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, VSP F1500, Virtual Storage Platform, Universal Storage Platform V/VM, or HUS VM storage systems for DP pools:
  - Parity groups from which volumes have already been created can be added
  - Drive type characteristics (for example: drive type and drive speed) in parity groups and RAID level
- Verify the following when using HUS 100 or AMS 2000 family of storage systems for DP pools:
  - Parity group capacity to add
  - Number of parity groups to add
  - There are drives in which a parity group has not been created

Note: In HDT pools, if different drive types and/or RAID levels are mixed in a single tier, they will all be considered equal for data placement regardless of page access frequency. As a result, I/O performance will be dependent on the drive type characteristics and RAID level on which any given page resides

In HDP pools, if different drive types and/or RAID levels are mixed in an HDP pool, I/O performance will be dependent on the drive type characteristics and RAID level on which any given page resides.

Procedure

1. On the **Resources** tab, expand the storage systems, select an HDP pool on the target storage system, and click **Expand Pool**.
2. To expand an HDP pool:
   a. In Additional Parity Groups, click **Add Parity Groups**.
   b. Select one of more parity groups, click **Add to Pool**, and then click **Close**. The Pool Summary is updated.

   Tip: To change an existing HDP pool into an HDT pool, click **Changes to an HDT pool**.

3. To expand an HDT pool, in Additional Parity Groups, choose **Standard** or **Mixed** mode.

   For **Standard** mode (single drive type/speed, chip type, RAID level, and volume location in a single tier):
   - Tier 1 is the default, click **Add Parity Groups**, select the parity group, click **Add to Pool**, and click **Close**. Select an available parity group that best meets your performance or capacity needs (Tier 1 for best performance, Tier 2 for next best performance, and Tier 3 for capacity).

- (Optional) Click **+** to add Tier 2 and Tier 3, configure the tiers based on your performance or capacity needs from volume choices in the **Add New Tier** dialog box. The Tier Configuration table in Pool Summary is updated.

Tip: To delete an existing tier, click **X** in the **Tier** tab.

For **Mixed** mode (mixed drive types/speeds, chip types, RAID levels, and/or volume locations in a single tier):

- Click **Add Parity Groups**.

Note: There are two parity group tab choices from which you can select: **Internal Parity Groups** and **External Parity Groups** (the **Internal Parity Groups** tab is set by default). If you select the **External Parity Groups** tab and select one or more parity groups, this enables the **External LDEV Tier Rank** menu from which you must choose a ranking for the tier.

- For mixed mode in the **Internal Parity Groups** tab or the **External Parity Groups** tab, select parity groups that you want to add to the HDT pool, click **Add to Pool**, and click **Close**. The Tier Configuration table in Pool Summary shows the new tier configuration status for each tier.

4. Click **Show Plan** and confirm that the information in the plan summary is correct. If changes are required, click **Back**.
5. (Optional) Update the task name and provide a description.
6. (Optional) Expand **Schedule** to specify the task schedule. You can schedule the task to run immediately or later. The default setting is **Now**. If scheduled for **Now**, select **View task status** to monitor the task after it is submitted.
7. Click **Submit**. If the task is to run immediately, the task begins.
8. You can check the progress and the result of the task on the **Tasks & Alerts** tab. Click on the task name to view details of the task.

Result

When the task is complete, the DP pools are expanded. You can verify the updated information on the DP Pools list.

Related concepts

- About virtualizing storage capacity on page 171

Related tasks

- Shrinking a DP pool on page 183
- Expanding DP volumes on page 185

- Reclaiming zero pages on page 186

Shrinking a DP pool

Shrinking a DP pool allows the recovery of excess free capacity.

Note: When you shrink a DP pool, Hitachi Command Suite automatically performs an immediate *quick* format when the volumes that were DP pool volumes are internal volumes, and a *basic* format when the volumes that were DP pool volumes are external volumes, but became basic volumes because those volumes were removed from the DP pool, which allows these basic volumes to be used again. A best practice is to run this task when system activity is low and major system operations are not running.

Procedure

1. On the **Resources** tab, select **Storage Systems** in the navigation tree.
2. Expand the tree for the target storage system that includes the DP pool you want to shrink, and select it in **DP Pools**.
3. In the **DP Pool Vols** tab, select DP pool volumes, and then click **Shrink Pool**.
4. Review the plan and specify additional information, as appropriate:
   - Verify the information that is displayed.
   - Enter a name in **Task Name**.
   - Specify when to execute the task.
5. Click **Submit**.

   The plan is registered as a task.
6. On the **Tasks & Alerts** tab, confirm that the task completed.
7. In the **Storage Systems** tree, return to the target storage system, click **DP Pools**, and confirm that the information is updated.

Result

Excess free capacity has been recovered from the DP pool.

Note: If a DP pool volume becomes blocked and the processing failed during execution of the Shrink Pool task, execute the `FormatLU` command in the Device Manager CLI. For details about the `FormatLU` command, see the *Hitachi Command Suite CLI Reference Guide*.

Related concepts

- About virtualizing storage capacity on page 171

Related tasks

- Expanding DP pools on page 180

Modifying DP pool settings

After setting up a DP pool, you can modify DP pool settings.

Procedure

1. On the **Resources** tab, select **DP Pools** under the target storage system.
2. From the list of DP pools, select the pool you want to modify and click **Edit Pool**.
3. Modify the settings as appropriate and click **Submit**.
4. On the **Tasks & Alerts** tab, confirm that all tasks are completed.
5. In the **Storage Systems** tree, return to the target storage system and click **DP Pools** to confirm that the information is updated.

Related concepts

- About virtualizing storage capacity on page 171

Related tasks

- Creating a DP pool on page 173

Deleting DP pools

After unallocating the DP volumes belonging to a DP pool, you can delete the DP pool.

Note: When you remove an encrypted DP pool in HUS 150 systems, this cancels the encryption for all drives that comprise the target DP pool, and releases the encryption on all DP pool volumes in the DP pool.

Before you begin

Before you delete DP pools, first unallocate the DP volumes that belong to the DP pools to be deleted.

Note: When you delete a DP pool, HCS automatically performs an immediate quick format only on those volumes that were DP pool volumes, but became basic volumes because the associated DP pools were deleted. This allows these basic volumes to be used again. A best practice for deleting a pool is to run this task when system activity is low and major system operations are not running.

Procedure

1. On the **Resources** tab, select **Storage Systems**.
2. Expand the tree, select the target storage system, and select **DP Pools**.
3. From the list of pools, select one or more target DP pools to delete, and click **Delete Pools**.

4. In the **Delete Pool** dialog box, confirm that the information matches the DP pools to be deleted.
   Optionally, update the task name and provide a description.
5. Click **Submit**.
   The delete DP pools process begins.
6. You can check the progress and the result of the delete DP pools task on the **Tasks & Alerts** tab
   Verify the results for each task by viewing the details of the task.

Result

The deleted DP pool no longer appears in the DP pool list for the target storage system.

Related concepts

- About virtualizing storage capacity on page 171

Related tasks

- Creating a DP pool on page 173

Related references

- Notes on performing quick formats on page 151

Expanding DP volumes

You can expand the size of a DP volume to expand its capacity.

Procedure

1. On the **Resources** tab, select **Storage Systems** in the navigation tree.
2. Expand the tree for the target storage system that includes the pool you want to modify, and select the DP pool in **DP Pools**.
3. On the **DP Vols** tab, select one or more volumes you want to expand and click **Expand DP Volume**.
4. Specify the new capacity for the volume.
5. Click **Submit**.
6. On the **Tasks & Alerts** tab, confirm that all tasks are completed.
7. In the **Storage Systems** tree, return to the target storage system, click the target DP pool and view the **DP Vols** tab to confirm that the information is updated.

Related concepts

- About virtualizing storage capacity on page 171

Related tasks

- Expanding DP pools on page 180

Reclaiming zero pages

Reclaiming unused zero pages for a DP pool releases unused capacity.

Procedure

1. On the **Resources** tab, select **Storage Systems** in the navigation tree.
2. Expand the tree for the target storage system that includes the appropriate DP pool, and select the DP pool in **DP Pools**.
3. On the **DP Vols** tab, select one or more volumes and click **Reclaim Zero Pages**.
4. Click **Submit**.
5. In the task list, confirm that the task is completed.

Related concepts

- About virtualizing storage capacity on page 171

Related tasks

- Expanding DP pools on page 180

Changing an access attribute to read/write

Before changing the access attribute from read-only or protect to read/write, considering the following:

- Do not assign an access attribute to a volume if any job is manipulating data on the volume. If you assign an access attribute to such a volume, the job will possibly end abnormally.
- Make sure that the retention term is expired. If expired, the Retention Term column in the **Data Retention** window shows 0. To change the access attribute to read/write within the retention term, contact customer support.
- Make sure that Expiration Lock indicates Disable > Enable. If it indicates Enable > Disable, changing to read/write is restricted by an administrator for some reason. Contact the administrator of your system to ask if you can change the access attribute. For details, see the *Provisioning Guide* for your storage system.

Procedure

1. On the **Resources** tab, click **Storage Systems**, and then expand **All Storage Systems** and the target storage system.
2. Choose one of the following options.
   - For VSP G series or VSP F series:
     Right-click the target storage system, and then select **Other Functions**.
   - For other available storage systems:

From the **Actions** list in the application pane, select **Element Manager**. Refer to the documentation for the native management tool for your storage system.

3. In Device Manager - Storage Navigator, click **Actions > Other Function > Data Retention** to open the **Data Retention** window.
4. Click  to change to **Modify** mode.
5. Select an LDKC number in the **LDKC** list, select a group in which the CU belongs in the **CU Group** list, then click a CU in the tree.
6. Right-click a volume for which you want to change access attributes. You may select multiple volumes, select **Attribute**, then click **Read/Write**.
7. Click **Apply** to apply the setting.

Related references

- Create Pool dialog box on page 176

Virtualizing storage tiers (HDT)

This module describes how to manage data relocation for HDT volumes.

About virtualizing storage tiers

Dynamic Tiering functionality lets you configure the monitoring frequency modes at which data is accessed and relocate that data in a specific pool and tier based on the results.

For example, you might create a pool that combines volumes having different cost performances, such as combining high-speed volumes (SSD, FMD, FMC, or SAS) with inexpensive low-speed volumes (SATA). The data in this pool is then automatically relocated among the volumes depending on the I/O load:

- High-load pages are allocated to high-speed volumes
- Low-load pages are allocated to low-speed volumes

By using Hitachi Command Suite, you can fine-tune settings related to the monitoring of HDT pools and data relocation depending on operating conditions.

You can configure a tiering policy in an HDT pool so that parity groups are assigned to specific tiers (Tier 1, Tier 2, and Tier 3) to take advantage of drive types, drive speeds, chip types, and RAID levels to optimize your performance and capacity needs. Tiering allows data in a host volume to be spread across multiple tiers (Tier 1 for high-speed, Tier 2 for next highest speed, and Tier 3 for additional capacity independent of drive type or RAID level), which provides the flexibility to specify settings that address your business conditions or the characteristics of certain applications. Tiering allows you to optimize data placement, improve volume performance, and reduce costs through more effective use of your resources.

HDT settings that can be specified using HCS include:

- Monitoring HDT pools and relocating data
  Monitoring and data relocation can be configured to run automatically by specifying a time or interval in advance, or you can manually monitor and relocate data as required. For VSP G1000, VSP G1500, VSP F1500, and VSP Gx00 models, data relocation speed can be set. For example, using a slower relocation speed will reduce the impact of relocation on other I/O requests.
- Specifying the buffer space for HDT pools
  When HDT pools are created or edited, on each hardware tier you can specify a ratio corresponding to the buffer space for new page assignment (an area reserved for increasing used capacity). Similarly, you can specify a ratio corresponding to the buffer space for tier relocation (a working area reserved for the storage system to use when relocating data). However, be aware that changing the default values might degrade performance.
- Applying a tiering policy and setting priority to HDT volumes
  You can regulate tiering to balance performance and cost, such as by preventing more data than is necessary from being placed in a fast hardware tier and thereby reducing management costs. After these settings are configured, regularly review the status of the settings of that tiering policy and the amount of capacity in use by each hardware tier to verify that resources are being appropriately allocated. Configuration settings can be changed any time if costs increase or performance is lower than expected.
  In an HDT pool, data with a high frequency of I/O operations is automatically preferentially placed in a high-speed hardware tier. Configure the following settings to control data placement, ensuring that important data is placed in a fast hardware tier, according to your business requirements.
  - Apply a tiering policy
    To ensure that data with high importance but with few I/O operations is placed in a hardware tier that maintains a certain speed at all times, you can specify the target hardware tier. You can also apply a tiering policy to HDT volumes that determines the capacity ratio of each hardware tier, by defining such a policy in advance.
  - Specify a new page assignment tier
    When HDT volumes are created or allocated, specify which hardware tier the new page of an HDT volume will be assigned with priority. Among the hardware tiers defined by a tiering policy, specify High for an upper-level hardware tier, Middle for a medium-level hardware tier, and Low for a low-level hardware tier.
  - Set relocation priority
    When HDT volumes are created or allocated, specify whether to prioritize relocation of the data of the target HDT volumes.
  - Real-time tier controls

If you use the Active Flash functionality, the page performance of tiered HDT volumes will be monitored, and data will be relocated to an appropriate tier if necessary due to changes in the I/O load across short intervals.

  - Optimizing tier allocation
    By using IOPH propagation, you can enable optimized tier allocation by transmitting performance monitoring data to HDT pools that contain volumes that are inaccessible from a host. These operations are performed from the CLI.

- Editing a tier rank for an external HDT pool volume
  When an external volume is included in an HDT pool volume that is one of the volumes making up an HDT pool, you can define the external LDEV tier rank (high, medium, or low) according to the performance.

A registered Tiered Storage Manager license enables the following:

- Evaluating and analyzing the operation status of HDT pools
  Use the Mobility tab to evaluate and analyze the operation status that is related to an HDT pool.
- Setting a schedule for relocating data and monitoring HDT pools
  Register the time of HDT pool monitoring and data replacement as a template schedule.
- Editing tier relocation for HDT volumes (preventing relocation by volume)
  Specify whether data can be relocated for each HDT volume. Tier relocation can be controlled according to the characteristics and operating status of applications using HDT volumes, such as by preventing other applications from relocating data of the volumes they are using when there is an application for which data relocation takes a high priority.
- Restoring a data placement by applying a data placement profile of HDT volumes
  Restore a previous data placement by saving data placements of optimized HDT volumes per page as profiles, and applying them according to operation. For example, if an HDT volume is being used for multiple operations that have different access characteristics (such as online and batch processing), you can create data placement profiles that fit the different processes and apply the appropriate profile before beginning processing. By doing so, you can restore a data placement that fits the characteristics of the target processing in advance, which prevents I/O performance from dropping. In addition, by setting up a schedule, you can update and apply profiles at regular intervals to suit the operation of applications. Profiles are applied only to pages placed in Tier 1 of HDT pools.



Tip: If you are using a VSP G1000, VSP G1500, or a VSP F1500 storage system, some windows might not display special HDT pool data, such as hardware tiers, for Dynamic Tiering for Mainframe volumes.

Related concepts

- Creating a storage operating environment on page 140

Related tasks

- Manually starting or stopping the monitoring of HDT pools on page 190
- Manually starting or stopping the tier relocation of an HDT pool on page 191
- Scheduling monitoring and tier relocation of HDT pools on page 191
- Editing tier relocation for HDT volumes on page 192
- Applying a tiering policy to HDT volumes on page 193
- Customizing a tiering policy for HDT volumes on page 193
- Creating a data placement profile for HDT volumes on page 195
- Updating a data placement profile for HDT volumes on page 196
- Editing a data placement profile for HDT volumes on page 197
- Applying a data placement profile for HDT volumes on page 198
- Scheduling data placement profiles for HDT volumes on page 199
- Editing an external LDEV tiering rank for an HDT pool on page 200

Related references

- Notes on data placement profiles for HDT volumes on page 194

Manually starting or stopping the monitoring of HDT pools

You can start or stop monitoring of an HDT pool manually.

Before you begin

A Tiered Storage Manager license is required to perform the operation from the Mobility tab.

After confirming that the HDT pools setting for Tier Management is set to Manual or Custom, you can start or stop monitoring of an HDT pool.

Procedure

1. From the tree view in the **Resources** tab, select **Storage Systems** (or use the **Mobility** tab).
2. Expand the tree and select **DP Pools** under the target storage system.
3. Select one or more HDT pools and click either the **Start Monitoring** or **Stop Monitoring** button, both found in the **More Actions** menu.
4. Set the desired items and execute the task.
5. View the list of tasks to check execution results.
6. Click the link for the task name and check that monitoring of each HDT pool has started or stopped.

Related concepts

- About virtualizing storage tiers on page 187

Manually starting or stopping the tier relocation of an HDT pool

You can manually start or stop tier relocation of an HDT pool.

Before you begin

A Tiered Storage Manager license is required to perform the operation from the Mobility tab.

After confirming the following, you can start relocation of an HDT pool:

- Existence of two or more hardware tiers in the target HDT pools
- The HDT pools setting for Tier Management is Manual or Custom

To stop relocation of an HDT pool, confirm:

- The HDT pool setting for Tier Management is Manual or Custom

Procedure

1. From the tree view in the **Resources** tab, select **Storage Systems** (or use the **Mobility** tab).
2. Expand the tree and select **DP Pools** under the target storage system.
3. Select one or more HDT pools and click either the **Start Relocation** or **Stop Relocation** button.
4. Set the desired items and execute the task.
5. View the list of tasks to make sure that all tasks have completed.
6. Click the link for the task name and check that tier relocation of each HDT pool has started or stopped.

Related concepts

- About virtualizing storage tiers on page 187

Scheduling monitoring and tier relocation of HDT pools

You can specify the schedule for monitoring and relocating HDT pools.

Before you begin

A Tiered Storage Manager license must be registered.

Procedure

1. From the **Mobility** tab, select **DP Pools**.
2. Expand the tree view and select the target storage system. From the list of DP pools, select the HDT pool for which the schedule is to be set.
3. On the **Summary** panel, click **Actions**, and select **Edit Pool**.
4. In the **Edit Pool** dialog box, click **HDT Settings**. In **Tier Management** select the **Custom** radio button. Click **Select Template** to use an

existing schedule template, or click **Create Template** to create a new schedule template.

5. Specify the necessary items by following the instructions in the window and create and execute the plan.
6. View a list of tasks to make sure that all tasks completed.
7. In the **Mobility** tab select the HDT pool. Then from **Tier Management** in **Summary** confirm the template names.

Related concepts

- About virtualizing storage tiers on page 187
- About data mobility on page 444

Editing tier relocation for HDT volumes

There are conditions where it is useful to disable tier relocation of volumes to prevent unnecessary data movement. For example, to preferentially perform tier relocation of volumes whose I/O activity varies greatly, and suppress tier relocation of other volumes.

Before you begin

A Tiered Storage Manager license must be registered.

Procedure

1. Click the **Mobility** tab, and select **Logical Groups** in the navigation pane.

   Logical Groups can contain both **Public Logical Groups** and **Private Logical Groups**. You can discover and register not only at the logical group level, but also at the host or HDT volume level.

   Note: If you want to register by host, you can start this process by selecting the **Resources** tab.

2. Expand the tree to select the target logical group from **Public Logical Groups** or **Private Logical Groups**.
3. If you want to set all HDT volumes in the logical group at the same time, select **Actions** located in the corner of the application pane, then **Edit Tier Relocation**. If you want to specify the target volumes, select one or more HDT volumes from the list of volumes, and then click the **Edit Tier Relocation** button.
4. Select **Enable** to enable tier relocation for HDT volumes. Select **Disable** to disable tier relocation.
5. Set any required items in the window, and then execute the task.
6. Verify that the task completed.

Result

The applied tier location settings can be checked from Summary or the list of volumes, which are displayed by selecting a logical group from the Mobility tab.

Related concepts

- About virtualizing storage tiers on page 187

Applying a tiering policy to HDT volumes

Before you begin

- Two or more hardware tiers must exist in an HDT pool.
- If selecting multiple HDT volumes from a logical group, all selected HDT volumes must belong to the same HDT pool.
- A Tiered Storage Manager license is required to perform operations from the Mobility tab.

Procedure

1. From the tree view in the **Resources** tab, select **Storage Systems**.
   You can also perform this step from the **Mobility** tab.
   If editing from logical groups, perform this step from the **Mobility** tab.
2. Expand the tree and select the target HDT pools.
3. Select one or more HDT volumes from the list of volumes on the **DP Vols** tab, and then click **Edit Tiering Policy**.
4. Select the tiering policy, and then execute the task.
   You can schedule the task to be executed later.
5. View the list of tasks to make sure that all tasks are complete.

Result

The tiering policy is applied.

Tip: The capacity ratio of Tier 1 might exceed the value specified for the maximum allocation threshold if a data placement profile and a tiering policy are being used concurrently.

Related concepts

- About virtualizing storage tiers on page 187

Customizing a tiering policy for HDT volumes

You can set the value for the allocation threshold of a hardware tier. When a tiering policy is specified for HDT volumes, make sure changing the definition will not cause a problem.

Before you begin

- A Tiered Storage Manager license must be registered to perform this task from the Mobility tab.
- Verify the tiering policy to be changed.

Procedure

1. On the **Resources** tab, expand the **Storage Systems** tree.
   Note that you can also update tiering policies on the **Mobility** tab.
2. For the applicable storage system, select **DP Pools**.
3. On the **Tiering Policies** tab, select a tiering policy to change, and click **Customize Tiering Policy**.
4. Specify the required items, and then submit the task.
5. After the task is complete, on the **Resources** tab, select **DP Pools**.
6. On the **Tiering Policies** tab, select the tiering policy you changed to verify that the can verify the changes.

Result

The tiering policy is updated.

Related concepts

- About virtualizing storage tiers on page 187

Notes on data placement profiles for HDT volumes

When using data placement profiles for HDT volumes, it is useful to understand the following:

- Data placement profiles are applied only to pages placed in Tier 1 of HDT pools.
- For pages to which both a data placement profile and tiering policy are set, the settings of the profile are prioritized.
- After a data placement profile has been applied and monitoring and data relocation of HDT pools has been performed, a data placement is restored. Therefore, users need to determine the timing at which to apply a profile by taking into account the period during which the monitoring and data relocation are performed.
- When performing a create, update, apply, or release of a data placement profile, or when searching for or repairing an inconsistency in a data placement profile, users cannot perform other operations for the target storage system.
- A task to create, update, apply, or release a data placement profile or to search for or repair an inconsistency in a data placement profile takes time to complete after it is executed. This time might affect other operations, so make a plan based on the displayed time estimate. Keep in mind that the larger the number of target volumes, the more likely that the estimated

time and the actual execution time will differ greatly. You can also stop tasks that are being executed, and then restart them.
- Regardless of the allocated resource groups, data placement profiles created by another user can also referenced in the list of profiles. However, operation that can be performed on such profiles are restricted based on the allocated resource group and role.
- If the number of managed data placement profiles exceeds 200, displaying the profile list might take some time. In this case, set the [Rows/page] of the profile list to no more than 100.
- When the operations listed below are performed, users need to perform a search inconsistency operation from the Refresh Storage Systems dialog because the profile that is actually applied to the volume in a storage system might not match the profile information that can be referenced from Hitachi Command Suite.
  - When the Hitachi Command Suite database was overwritten, due to an import or restoration, while the applied profile existed.
  - When the storage system that contains an applied profile is deleted, and the storage system is re-registered.

Tip: If inconsistencies are detected, a message will be displayed in Manage Data Placement Profiles dialog. From the link shown in the message, perform the `repair inconsistencies` operation and release the profile. This can avoid unintended data to be fixed into a high-speed hardware tier by a profile that was applied in the past.

Related concepts

- About virtualizing storage tiers on page 187

Creating a data placement profile for HDT volumes

To restore a data replacement appropriate for application processing, the user creates a data placement profile when HDT volumes provide sufficient performance. A profile can be created for each logical group, and is managed with the corresponding logical group.

Before you begin

- A Tiered Storage Manager license must be registered
- Gather the name of the target logical group
- Specify settings for checking performance information, such as settings for linking to Hitachi Tuning Manager, or settings for the performance monitoring software of each application.

Tip: If linked to Hitachi Tuning Manager, you can view performance trend charts for the volumes in the logical group from the **Manage Data Placement Profiles** dialog box. For details on the linking method, see the *Hitachi Command Suite Administrator Guide*.

Procedure

1. From the **Mobility** tab, **General Tasks** pane, select **Manage Data Placement Profiles**.
2. Click the **Logical Group View** button, and specify the necessary items. Check the performance of the target logical group by using the performance trend charts or software for checking performance information, and then click the **Create Profile** button.
3. To create a plan, specify the necessary items by following the instructions in the window.
4. If there is no problem with the plan, execute it.
5. View the list of tasks to see the execution results.

Result

The created profile can be checked in the list of profiles displayed by clicking the Logical Group View button in the Manage Data Placement Profiles dialog box.

Tip: To delete data placement profiles, select the data placement profiles to be deleted, and click Delete Profiles.

Note: Regardless of the allocated resource groups, data placement profiles created by another user can also be referenced in the list of profiles. However, operations that can be performed on such profiles are restricted based on the allocated resource group and role. Therefore, specify profile names and descriptions to make it clear which operations are available.

Related concepts

- About virtualizing storage tiers on page 187

Updating a data placement profile for HDT volumes

You can update an existing data placement profile to reflect the latest HDT volume data placement in the profile.

Tip: To change the target HDT volumes of a profile due to changes in the logical group configuration or other reasons, re-create the profile.

Before you begin

- Register a Tiered Storage Manager license.
- Determine the name of the target logical group.
- Specify settings for checking performance information, such as settings for linking to Hitachi Tuning Manager, or settings for the performance monitoring software of each application.

Tip: If linked to Hitachi Tuning Manager, you can view performance trend charts for the volumes in the logical group from the **Manage Data Placement Profiles** dialog box. For details on the linking method, see the *Hitachi Command Suite Administrator Guide*.

Procedure

1. From the **Mobility** tab, **General Tasks** pane, select **Manage Data Placement Profiles**.
2. Click the **Logical Group View** button, specify the necessary items. Check the performance of the target logical group by using the performance trend charts or software for checking performance information, select the row of the profile to be updated (only one profile can be selected) from the list of data placement profiles, and then click the **Update Profile** button.
3. To create a plan, specify the necessary items by following the instructions in the window.
4. If there is no problem with the plan, execute it.
5. View the list of tasks to see the execution results.

Result

The updated profile information can be viewed from the profile operation history list, which is displayed by clicking the Logical Group View button in the Manage Data Placement Profiles dialog box.

Tip: When you want to periodically update the profile according to the intervals for monitoring HDT pools and replacing data, you can set a schedule for updates by using the Schedule Profiles button. Also, When an applied profile is updated, the profile in a condition before the update will continue to be used until the updated profile is reapplied.

Related concepts

- About virtualizing storage tiers on page 187

Editing a data placement profile for HDT volumes

You can change the name and description of an existing data placement profile.

Before you begin

A Tiered Storage Manager license must be registered.

Procedure

1. From the **Mobility** tab, **General Tasks** pane, select **Manage Data Placement Profiles**.

2. Click **Overall Profile View** or **Logical Group View**, select the row of the profile to be edited from the list of data placement profiles (only one profile can be selected), and then click **Edit Profile**.
3. Edit the necessary items and submit the task.

Result

When the task completes, the data placement profile is updated.

Related concepts

- About virtualizing storage tiers on page 187

Applying a data placement profile for HDT volumes

Before beginning application processing, apply a data placement profile to restore a data placement that fits the characteristics of the processing. After confirming that the data placement is restored and performance is improved, release the applied profile to return to normal HDT operation.

Tip:
- To apply or release the profile periodically to match the application operations, click Schedule Profiles to schedule when to apply and release the profile. When you apply the profile, you can also specify a release schedule.
- The capacity ratio of Tier 1 might exceed the value specified for the maximum allocation threshold if a data placement profile and a tiering policy are being used concurrently.

Before you begin

- A Tiered Storage Manager license must be registered
- A data placement profile must be created
- Gather the name of the target logical group
- Specify settings for checking performance information, such as settings for linking to Hitachi Tuning Manager, or settings for the performance monitoring software of each application

Tip: If linked to Hitachi Tuning Manager, you can view performance trend charts for the volumes in the logical group from the **Manage Data Placement Profiles** dialog box. For details on the linking method, see the *Hitachi Command Suite Administrator Guide*.

Procedure

1. On the **Mobility** tab, **General Tasks** pane, select **Manage Data Placement Profiles**.
2. Click the **Overall Profile View** button, or **Logical Group View** button, to view the creation date, usage, and effect of applying profiles in the past, and select the row of the profile to be applied (only one profile can

be selected) from the list of data placement profiles, and then click the **Apply Profile** button.

3. To create a plan, specify the necessary items by following the instructions in the window.
4. If there is no problem with the plan, execute it.
5. View the list of tasks to see the execution results.

   After the monitoring and data relocation of HDT pools finish, perform the following operations to check the effect of applying the profile, and to release the applied profile.
6. Using the software for checking the performance information, check the effects of applying the profiles.

   If linked with Hitachi Tuning Manager, from the **Manage Data Placement Profiles** dialog box, click **Logical Group View** to check the performance trend chart of the target logical group.
7. To return to normal HDT pools operation, click **Release Profile**, specify the necessary items, and then release the applied profile.

Related concepts

- About virtualizing storage tiers on page 187

Scheduling data placement profiles for HDT volumes

You can set weekly or monthly schedules for applying, releasing, and updating data placement profiles.

Tip: If multiple operation schedules are registered for a single logical group, the task list displays the operation that will be executed first. For example, if a schedule to apply a profile is set followed by a schedule to release the profile, then the application task is displayed until the profile is applied. After the profile is applied, the task to release the profile is displayed.

Before you begin

- Register a Tiered Storage Manager license.
- Create a data placement profile.
- Identify the name of the target logical group.

Procedure

1. On the **Mobility** tab, select **General Tasks**, then select **Manage Data Placement Profiles**.
2. Click the **Overall Profile View** button, or **Logical Group View** button, then click the **Schedule Profiles** button.
3. To create a plan, specify the necessary items by following the instructions in the window.
4. If there is no problem with the plan, submit it.
5. View the list of tasks to make sure that the operation for which a schedule is set is registered as a task.

Related concepts

- About virtualizing storage tiers on page 187

Editing an external LDEV tiering rank for an HDT pool

You can edit the external LDEV tiering rank (Low/Medium/High) for a pool volume.

Before you begin

- Virtualize sufficient storage.
- Define hardware tiers that consist of external volumes.
- Connect the external storage system.

Procedure

1. On the **Resources** tab, select **Storage Systems**, and then select the target **DP Pools**.
2. Select the target HDT pool, and click the **DP Pool Vols** tab.
3. Select the target HDT pool volume the **DP Pools Vols** list, and then click **Edit External LDEV Tier Rank**.
4. In the **External LDEV Tier Rank** menu, change the tier ranking (Low/Middle/High) to a different value.
5. (Optional) Update the task name and provide a description.
6. (Optional) Expand **Schedule** to specify the task schedule. You can schedule the task to run immediately or later, The default setting is **Now**. If scheduled for **Now**, select **View task status** to monitor the task after it is submitted.
7. Click **Submit**. If the task is to run immediately, the task begins.
8. You can check the progress and the result of the task on the **Tasks & Alerts** tab. Click on the task name to view details of the task.

Result

The external LDEV tiering rank for the HDT pool volume has been changed on the target storage system.

Related concepts

- About virtualizing storage tiers on page 187

Allocating and unallocating volumes

This module describes volume allocation, provides information about volume allocation requirements, describes multiple ways in which you can allocate volumes, and provides related procedures for allocating volumes and editing volume or host access.

About allocating volumes

Volume allocation makes storage capacity available to host applications and file servers. Hosts and file servers must already be registered before volume allocation.

Depending on your registered storage systems, volumes can be allocated using basic volumes, pool volumes, or volumes from a tier. Basic volumes are volumes from a parity group. Any storage system can provide basic volumes. Allocating DP pool volumes involves grouping resources on storage systems that support this feature. DP pools must exist prior to volume allocation from a DP pool. To allocate volumes from a tier, a Tiered Storage Manager license is required, and tier policies must be established for storage systems that support tier policy configuration.

The following rules and behaviors apply when allocating volumes:

- Settings that are assigned when allocating volumes to a host become the default settings for the next time you allocate a volume to the same host. You can change these settings during volume allocation.
- If a variety of volumes with different characteristics have been allocated to a host, when you allocate a new volume, you can select an existing volume to set the volume allocation criteria.
- When allocating a volume, if no volumes match the specified requirements, new volumes are created from unused capacity and allocated to the host. When a basic volume is created, the volume is also formatted at the same time.
- You can use keyword or criteria-based searches to find existing unallocated volumes that meet your requirements.
- When you allocate volumes to a host, LUN paths are assigned automatically.
  LUN paths are storage port to host port mappings that provide host access to volumes. The host port is a host bus adapter (HBA) or an iSCSI name. WWN nicknames can be displayed to confirm target HBAs for a host while editing LUN paths during, or after, volume allocation.
- Volumes can be allocated on ports where LUN security is not enabled. All hosts with access to the port can access the volume.
- Volume allocation to a clustered host should be done by using logical groups. Using logical groups ensures that the same volumes are allocated to all hosts in the cluster.
- Volume allocation is not complete until you approve the volume allocation plan and submit the volume allocation task.
- When volumes are allocated to a host OS, you must create a file system on them and mount them, before they can be used by host applications.
- When you are linking with NAS Platform v11.3 or later and using VSP G400, G600, G800, and VSP F400, F600, F800 with NAS modules, and volumes are allocated for creating or expanding storage pools, it is recommended that you allocate volumes using the Create Storage Pool or

Expand Storage Pool dialog boxes. Device Manager can automatically specify a volume and path configuration for allocating volumes that follows the best practices for configuring storage pools.
- When volumes are allocated to a NAS Platform F or Data Ingestor, and Hitachi File Services Manager v3.2 is installed, you can create a file system and allocate the volume by using the Create File System dialog box.

Note: Before allocating volumes, review the available volume allocation methods. Understanding the available starting points for volume allocation will enable you to perform volume allocation in a way that best suits your requirements, and will help you understand the Allocate Volumes dialog box and the Define Clustered-Host Storage dialog box.

Related concepts

- Allocating storage on page 141

Related tasks

- Allocating volumes to selected file servers on page 206

Related references

- Prerequisites for allocating volumes on page 204
- Notes on performing quick formats on page 151
- Volume allocation methods on page 202
- Allocate Volumes dialog box on page 213
- Workflow for providing NAS Platform file shares on page 250
- Notes on setting LUN paths for NAS Platform on page 257

Volume allocation methods

When you allocate volumes, you can use the Allocate Volumes dialog box or the Define Clustered-Host Storage dialog box.

Begin by selecting resources, such as volumes or hosts, and then click Allocate Volumes, which opens the Allocate Volumes dialog box.

Click Allocate Volumes without first selecting a resource, such as a volume or host.

There are multiple methods in which to begin allocating volumes. The visible fields in the Allocate Volumes dialog box will vary depending on the starting point and whether you have selected a resource to begin the procedure. For example, if you select resources first, the Allocate Volumes dialog box will prompt you for less information.

If you are linking with NAS Platform v11.3 or later and allocating volumes to a file server when creating or expanding storage pools, the following is recommended:

- Use the Create Storage Pool dialog box or the Expand Storage Pool dialog box instead of the Allocate Volumes dialog box to create or expand storage pools and allocate volumes at the same time.

Device Manager can automatically specify a volume and path configuration for allocating volumes that follows the best practices for configuring storage pools.

When using the Allocate Volumes dialog box, the following conditions apply when you begin volume allocation using these available methods and starting points:

- From the General tasks pane
  On the Resources or Mobility tab, from General Tasks, click Allocate Volumes. Because no resource was selected first, you must specify host, volume, and other volume allocation requirements.
- Selecting hosts or file servers
  From the Resources tab, select one or more hosts or file servers, and click Allocate Volumes. The dialog will display the selected host or file server names instead of prompting.
- Selecting volumes
  From the Resources tab, select one or more volumes, and click Allocate Volumes. The dialog will prompt for host or file server name, but will not prompt for volume criteria.
- Selecting clustered hosts
  Configuring clustered hosts in a logical group helps ensure that volume allocations will be consistent for all members of the cluster. From the Resources tab, select a logical group, select all host members of the group, and then click Allocate Volumes. If clustered hosts are not in a logical group, from General Tasks or Actions, click Allocate Volumes and select multiple hosts. You may also use the Resources tab to locate clustered hosts, and then Allocate Volumes. The use of logical groups is recommended to avoid host selection errors.
- Searching for volumes
  You can search for volumes that meet specific criteria (such as storage system, specific allocation status, volume type, drive type, drive speed, or chip type), select one or more volumes from the search results, and click Allocate Volumes.
  You can also search using full or partial text keywords such as storage system name or host name, and then allocate volumes to the selected resource. Searching eliminates the need to manually navigate to the resource.
- Using existing volume settings
  From the Resources tab, select an allocated volume that has desired attributes and criteria for the new volume allocation, and click Allocate Like Volumes.

When you use the Define Clustered-Host Storage dialog box for Fibre Channel and Fibre Channel over Ethernet (FCoE) connections to create a cluster configuration using existing hosts, or to add a new host to an existing host group or cluster, ensure that you allocate the same volumes that are assigned to the existing hosts in the host group to the new host.

Related concepts

- About allocating volumes on page 201

Related tasks

- Allocating volumes from general tasks on page 204
- Allocating volumes to selected hosts on page 205
- Allocating volumes to selected file servers on page 206
- Allocating selected volumes to hosts on page 208
- Allocating volumes to clustered hosts on page 209
- Allocating volumes by using a keyword search on page 210
- Allocating volumes by using a criteria search on page 211
- Allocating volumes by using existing volume settings on page 212

Related references

- Allocate Volumes dialog box on page 213

Prerequisites for allocating volumes

Before allocating volumes to a host or file server, you must verify that the host, file server, and storage systems are registered in Hitachi Command Suite.

If you want to select volumes from a tier, Tiered Storage Manager license must be enabled.

In addition, determine:

- The target host or file server
- The volume type, count, capacity, and performance characteristics
- An appropriate storage system

Related concepts

- About allocating volumes on page 201

Allocating volumes from general tasks

The General Tasks list is conveniently located for volume allocation (and other frequent tasks). Because you have not selected a specific resource (for example, a volume or a host) the dialog will prompt for all necessary information.

Procedure

1. On the **Resources** or **Mobility** tab, select **Allocate Volumes**.

   Tip: If you do not see **Allocate Volumes** listed, click **more...** to see all menu items.

2. On the **Allocate Volumes** dialog box, specify volume allocation requirements.
3. Click **Show Plan** and confirm that the information in the plan summary is correct. If changes are required, click **Back**.
4. (Optional) Update the task name and provide a description.
5. (Optional) Expand **Schedule** to specify the task schedule.

   You can schedule the task to run immediately or later. The default setting is **Now**. If the task is scheduled to run immediately, you can select **View task status** to monitor the task after it is submitted.
6. Click **Submit.**
   If the task is scheduled to run immediately, the process begins.
7. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.

Result

A completed task indicates a successful volume allocation.

Related references

- Prerequisites for allocating volumes on page 204
- Volume allocation methods on page 202
- Allocate Volumes dialog box on page 213

Allocating volumes to selected hosts

When one or more hosts are selected, the hosts are displayed in the **Allocate Volumes** dialog box. The dialog box prompts for all additional information.

Procedure

1. On the **Resources** tab, select **Hosts**.
2. Hosts are grouped by operating system. Click the operating system for the target hosts.

   Tip: If you do not know the host OS, searching for the host may be the better method for locating the host and allocating volumes.

3. Select one or more hosts, and click **Allocate Volumes**.

4. In the **Allocate Volumes** dialog box, specify volume allocation requirements.
5. Click **Show Plan** and confirm that the information in the plan summary is correct. If changes are required, click **Back**.
6. (Optional) Update the task name and provide a description.
7. (Optional) Expand **Schedule** to specify the task schedule.

   You can schedule the task to run immediately or later. The default setting is **Now**. If the task is scheduled to run immediately, you can select **View task status** to monitor the task after it is submitted.
8. Click **Submit.**
   If the task is scheduled to run immediately, the process begins.
9. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.

Result

A completed task indicates a successful volume allocation.

Related tasks

- Viewing HCS task status on page 507

Related references

- Prerequisites for allocating volumes on page 204
- Volume allocation methods on page 202
- Allocate Volumes dialog box on page 213

Allocating volumes to selected file servers

When a file server or file server cluster are selected, the file server is displayed in the volume allocation dialog. The dialog will prompt for all additional information.

To allocate volumes to a file server or file server cluster:

Procedure

1. On the **Resources** tab, select **File Servers**, then select **All File Servers**.
2. In the Servers/Clusters list, select the row of the target file server or cluster (only one row can be selected), and click **Allocate Volumes**.
   The Allocate Volumes dialog box will launch.
3. Specify your volume allocation requirements.
4. Click **Show Plan** and confirm that the information in the plan summary is correct. If changes are required, click **Back**.
5. (Optional) Update the task name and provide a description.
6. (Optional) Expand **Schedule** to specify the task schedule.

You can schedule the task to run immediately or later. The default setting is **Now**. If the task is scheduled to run immediately, you can select **View task status** to monitor the task after it is submitted.

7. Click **Submit.**
   If the task is scheduled to run immediately, the process begins.
8. You can check the progress and the result of the task on the **Tasks & Alerts** tab. Click on the task name to view details of the task.

Tip: If you are linking with NAS Platform v11.3 or later and allocating volumes to a file server when creating or expanding storage pools, the following is recommended:

Use the **Create Storage Pool** dialog box or the **Expand Storage Pool** dialog box instead of the **Allocate Volumes** dialog box to create or expand storage pools and allocate volumes at the same time.

Device Manager can automatically specify a volume and path configuration for allocating volumes that follows the best practices for configuring storage pools.

Tip: For NAS Platform family, in cases such as node additions, if you want to allocate volumes to file server nodes that are in clusters, use the **Define Clustered-Host Storage** dialog box.

A completed task indicates a successful volume allocation.

9. To verify the volume allocation, select the file server or cluster to which the volumes were allocated, and then select the **System Drives** tab for NAS Platform family, or select the **Volumes** tab for NAS Platform F or Data Ingestor to confirm that the volumes were allocated successfully.

Result

Volumes have been allocated to a file server, or file server cluster, and verified.

Related concepts

- About registering and removing file servers on page 81
- About allocating volumes on page 201

Related references

- Prerequisites for allocating volumes on page 204
- Volume allocation methods on page 202
- Allocate Volumes dialog box on page 213
- Workflow for providing NAS Platform file shares on page 250
- Workflow for providing HNAS F and HDI file shares (HFSM) on page 275

Allocating selected volumes to hosts

When volumes are selected, they are displayed in the **Allocate Volumes** dialog box, which prompts you for all additional information to allocate selected volumes to one or more hosts.

Procedure

1. On the **Resources** tab, select **Storage Systems**.
2. Expand **All Storage Systems**, and select a specific storage system. Links for volume groups are listed in the **Detailed Information** list on the application pane. Listed groups represent available pools, parity groups, and volumes.
3. Select a group from the **Resources** tree, or click a link from the **Detailed Information** list.
   Individual pools, parity groups, and volumes are listed.
4. Select an appropriate resource. For example, select a pool or parity group with available capacity, or select an open-unallocated volume.
5. Click **Allocate Volumes**.

   Tip: Open-reserved volumes cannot be allocated. Open-allocated can be used to allocate like volumes.

6. On the **Allocate Volumes** dialog box, specify volume allocation requirements.
7. Click **Show Plan** and confirm that the information in the plan summary is correct. If changes are required, click **Back**.
8. (Optional) Update the task name and provide a description.
9. (Optional) Expand **Schedule** to specify the task schedule.
   You can schedule the task to run immediately or later. The default setting is **Now**. If the task is scheduled to run immediately, you can select **View task status** to monitor the task after it is submitted.
10. Click **Submit.**
    If the task is scheduled to run immediately, the process begins.
11. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.

Result

A completed task indicates a successful volume allocation.

Related references

- Prerequisites for allocating volumes on page 204
- Volume allocation methods on page 202
- Allocate Volumes dialog box on page 213

Allocating volumes to clustered hosts

To prevent possible errors with cluster host selection, group hosts that belong to a cluster into a public logical group or a private logical group. Alternatively, you can select multiple hosts or file servers in the **Allocate Volumes** dialog box.

Tip: When you create a cluster configuration using existing hosts or when you add a new host to an existing host group or cluster, use the **Define Clustered-Host Storage** dialog box to allocate volumes to the new host. When using the **Define Clustered-Host Storage** dialog box, you can add the WWN of the new host to the same host group as the WWN of the existing hosts.

Before you begin

- Add clustered hosts into the desired logical group type.
- Confirm the names of the hosts that belong to the target cluster (if you have not already added clustered hosts into a logical group).

Procedure

1. On the **Resources** tab, select **Logical Groups**.
2. Expand the **Public Logical Groups** or **Private Logical Groups** root folder, and locate an appropriate logical group.
3. Locate the **Hosts** and **Volumes** tabs under the summary pane, and select all hosts for the cluster.

   Note: You can allocate volumes to cluster hosts that are not in a logical group. However, you must ensure that your host selection is correct. Logical groups are strongly recommended.

4. Click **Allocate Volumes**.

   Tip: To allocate volumes to cluster hosts, verify that all cluster host names are displayed in the **Allocate Volumes** dialog box to prevent an incorrect allocation.

5. On the **Allocate Volumes** dialog box, specify volume allocation requirements.
6. Click **Show Plan** and confirm that the information in the plan summary is correct. If changes are required, click **Back**.
7. (Optional) Update the task name and provide a description.
8. (Optional) Expand **Schedule** to specify the task schedule.

You can schedule the task to run immediately or later. The default setting is **Now**. If the task is scheduled to run immediately, you can select **View task status** to monitor the task after it is submitted.

9. Click **Submit.**
   If the task is scheduled to run immediately, the process begins.
10. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.

Result

A completed task indicates a successful volume allocation.

The same volumes have been allocated to the individual hosts that belong to a cluster.

Related tasks

- Viewing HCS task status on page 507

Related references

- Prerequisites for allocating volumes on page 204
- Volume allocation methods on page 202
- Allocate Volumes dialog box on page 213

Allocating volumes by using a keyword search

A keyword search is a full or partial text search. For example, you can search for storage systems, hosts, volumes, parity groups, DP pools, and logical groups by entering a full or partial name. Note that file servers are not subject to search. Using a keyword search provides an alternate method of navigating to a resource to allocate volumes.

Procedure

1. Enter a value in the search box and press **Enter** on your keyboard. **All Resources** is the default selection criteria. You can limit the scope of searched resources by using the drop-down menu.

   Note: As the number of searchable objects grows, identify a specific resource to search, for example **Logical Groups**, to decrease search time.

2. Click the link for appropriate resource. For example, assuming you searched for hosts, from your search results, click the host name link. The appropriate location in the **Resources** tab is displayed. The host is effectively selected.
3. Click **Actions > Allocate Volumes**.
4. On the **Allocate Volumes** dialog box, specify your volume allocation requirements.

5. Click **Show Plan** and confirm that the information in the plan summary is correct. If changes are required, click **Back**.
6. (Optional) Update the task name and provide a description.
7. (Optional) Expand **Schedule** to specify the task schedule.

   You can schedule the task to run immediately or later. The default setting is **Now**. If the task is scheduled to run immediately, you can select **View task status** to monitor the task after it is submitted.
8. Click **Submit.**
   If the task is scheduled to run immediately, the process begins.
9. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.

Result

A completed task indicates a successful volume allocation.

Related references

- Prerequisites for allocating volumes on page 204
- Volume allocation methods on page 202
- Allocate Volumes dialog box on page 213

Allocating volumes by using a criteria search

Using a criteria search allows finding volumes with specific attributes, such as drive type or RAID level, and allocating them.

Procedure

1. From the search box drop-down menu, select **More Searches**.
2. Specify the search criteria and execute your search. There are **Basic** and **Advanced** tabs presenting criteria options. Basic criteria for volume status, type, drive performance, RAID level, or capacity requirement should meet most needs. Advanced criteria are more specialized.

   Tip: Searches can be saved for re-use, and can be saved as private or public. Note that saving or sharing a search requires the Tiered Storage Manager license.
3. From the search results, select the volumes and click **Allocate Volumes**.
4. On the **Allocate Volumes** dialog box, specify your host, and other volume allocation requirements.
5. Click **Show Plan** and confirm that the information in the plan summary is correct. If changes are required, click **Back**.
6. (Optional) Update the task name and provide a description.
7. (Optional) Expand **Schedule** to specify the task schedule.

You can schedule the task to run immediately or later. The default setting is **Now**. If the task is scheduled to run immediately, you can select **View task status** to monitor the task after it is submitted.

8. Click **Submit.**
   If the task is scheduled to run immediately, the process begins.
9. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.

Result

A completed task indicates a successful volume allocation.

Related references

- Prerequisites for allocating volumes on page 204
- Volume allocation methods on page 202
- Allocate Volumes dialog box on page 213

Allocating volumes by using existing volume settings

If you have allocated at least one volume to a host, you can allocate new volumes for the host by using the attributes of an existing volume as the default for the Allocate Volume dialog box. Your volume allocation can be exactly like, or similar to, the selected existing volume.

Procedure

1. On the **Resources** tab, select **Hosts** or **Logical Groups**.
2. Expand the tree for host or logical group type (public logical groups or private logical groups), and select the desired host or logical group to display a list of existing volumes.
3. Select an existing volume that matches or is similar to your requirement, and click **Allocate Like Volumes**.
4. The selected host or logical group name displays, and you are prompted for other volume allocation requirements.
   You can retain the existing settings, or change them as needed.
5. Click **Show Plan** and confirm that the information in the plan summary is correct. If changes are required, click **Back**.
6. (Optional) Update the task name and provide a description.
7. (Optional) Expand **Schedule** to specify the task schedule.
   You can schedule the task to run immediately or later. The default setting is **Now**. If the task is scheduled to run immediately, you can select **View task status** to monitor the task after it is submitted.
8. Click **Submit.**
   If the task is scheduled to run immediately, the process begins.
9. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.

Result

A completed task indicates a successful volume allocation.

Related references

- Prerequisites for allocating volumes on page 204
- Volume allocation methods on page 202
- Allocate Volumes dialog box on page 213

Allocate Volumes dialog box

Allocating volumes is the process for making storage capacity available to host applications and file servers. Hosts and file servers must already be registered prior to volume allocation.

When you enter the minimum required information in this dialog box, the Show Plan button activates to allow you to review the plan. Click the Back button to modify the plan to meet your requirements.

The following table describes the dialog box fields, subfields, and field groups. A *field group* is a collection of fields that are related to a specific action or configuration. You can minimize and expand field groups by clicking the double-arrow symbol (>>).

As you enter information in a dialog box, if the information is incorrect, errors that include a description of the problem appear at the top of the box.

Table 13 Allocate Volumes dialog box

Field	Subfield	Description
Host	-	If you select one or more hosts (including virtualization servers, or file servers) prior to clicking Allocate Volumes, the host names are displayed. If you did not select hosts, you are prompted to do so. The host drop-down list displays a single host, by name. Additionally, Select Hosts allows you to select multiple hosts, and displays related details such as the OS and WWN. For example, you can select cluster hosts for volume allocation. The host drop-down list provides a special selection called Allocate without LUN security. When you enter adequate volume information, LUN Path Options expand to a list of storage ports with LUN security disabled. Select one or more ports, and click Add to add them to the Selected Storage Ports list as part of this dialog box. All hosts with LUN paths on the selected storage port have access to the volumes on the ports.
Allocation Type	-	To allocate global-active device pairs, select Global-Active Device, which will display Primary/Secondary tabs for configuring paired volumes. For all other volume allocations, select Standard. Global-Active Device is not displayed if a selected host is a file-server.
No. of Volumes	-	Specify the number of volumes to allocate to the selected hosts.

Field	Subfield	Description
		For global-active device, this is the number of volumes to become paired volumes and allocated to the selected hosts.
Volume Capacity	-	Specify the volume size and select the unit of measure (for example, GB represents Gigabytes). The total capacity to be allocated (No. of Volumes * Capacity) is displayed. **Note:** You can create a new volume when volume criteria, such as volume capacity, cannot be satisfied by existing volumes. This is indicated in Advanced Options > Creating Volume Settings > Advanced Options by formatting and stripe size options for the new volume.
Storage system	-	If you select a storage system prior to clicking Allocate Volumes, its name is displayed. Otherwise, select the storage system that is providing the volumes. Any storage system can provide basic volumes. Volumes to be allocated from DP pools must be provided by a storage system that supports this feature. Allocating volumes from a tier requires configured tier policies for the storage system. The Any option means that the volume can be sourced from any storage system that meets the criteria of the dialog box. For global-active device, from the primary/secondary tabs, you must select primary and secondary storage systems that have been previously identified in the global-active device setup.
Virtual Storage Machine	-	This displays a list of one or more virtual storage machines from which you can choose during volume allocation. The virtual storage machines can also be viewed on the Administration tab by selecting Virtual Storage Machine. When allocating global-active device paired volumes, from the primary/secondary tabs, you must select the virtual storage machine to be used by both the primary and secondary storage for access to virtualized resources. If the virtual storage machine does not exist, it can be created during the global-active device setup.
Volume type	-	Select the volume type (for example, Basic Volume, Dynamic Provisioning, or Dynamic Tiering). The displayed volume types are determined based on your selected storage system. If you do not see an expected volume type, check that you have selected the correct storage system. The selected volume type affects the Advanced Options fields described below. For example, selecting a Basic Volume will populate Advanced Options > Volume Selection > Automatic with a default drive type, drive speed, chip type, RAID level, and parity group. These specifications may be altered. For example, if you change the drive speed, the parity group may change automatically. These changes are determined by available resources. You may also select an available tier from which to allocate the volume, or manually locate available unallocated volumes.

Field	Subfield	Description
		If you select volume type Dynamic Provisioning, Advanced Options > Volume Selection > Automatic, this populates with an HDP pool instead of a parity group (and the specifications may be altered). If you select Dynamic Tiering, the Volume Selection field displays. Volume selection is either automatic or manual (see the following section on the volume selection field). The Any option means that the volume can be sourced from any available volume type that meets the capacity criteria, and is available from the storage system. For NAS Platform versions earlier than 12.1, basic volumes are displayed by default. If either HDP (Dynamic Provisioning) or HDT (Dynamic Tiering) are selected, you must manually select the volume to assign (see the following Volume Selection, Manual option).
Volume location	-	Select Internal if you want your volume to come directly from your selected storage system, or External if volumes are mapped from external storage systems (virtualized volumes) and are adequate for your needs. The external option can appear for basic volumes and HDP pools. When viewing a list of DP pools, the volume location column indicates whether it is internal or external. **Note:** This field is replaced by Volume Selection when Dynamic Tiering is selected (see the previous Volume type section).
Volume Selection	-	This field displays only if Dynamic Tiering is selected (see volume type above). **Note:** This field is not the same field as the Volume Selection field defined in the following Advanced Options section. Select Automatic to allow the system to configure volume allocation space for you from available DP pools. Select Automatic and click Select Pool to list the available DP pools. Make an appropriate selection that corresponds to your capacity requirements, and click OK. Verify that the name is correctly displayed. When displayed under Advanced Options, you can tailor the Tiering Policy Setting (see the following section). When you submit the task, a volume is created for the host. Select Manual to display the selected volumes list, select Add Volumes,and select Unallocated Volumes or Allocated Volumes. Choose an appropriate volume, click Add to update Selected Volumes, and click OK to return to volume allocation.
>> Advanced Options	(See the following twelve fields for details)	When you set the volume type to Basic Volume or Dynamic Provisioning, this displays the Tier, Automatic, and Manual options. When the volume type is Dynamic Tiering, the fields display as explained in Volume Type. The fields can be displayed or hidden when you click Advanced Options. These fields support explicit volume allocation (Tier or Manual) or volume allocation based on criteria such as drive type, drive speed, chip type, or RAID level.

Field	Subfield	Description
Tiering Policy Setting	-	Displays only if Dynamic Tiering is selected as the volume type, and an HDT pool has been selected with Select Pool (see previous Volume Selection section). You can select a specific tier policy for the volume to be allocated, or select All.
New Page Assignment Tier	-	For VSP Gx00 models, VSP G1000, VSP G1500, VSP F1500, VSP, or HUS VM storage systems, selecting this option specifies to which hardware tier the new page of an HDT volume is to be assigned with a specified priority. Within the hardware tiers for which the tiering policy is set, specify *High* for an upper-level hardware tier, *Middle* for a medium-level hardware tier, and *Low* for a low-level hardware tier.
Relocation Priority	-	For VSP Gx00 models, VSP G1000, VSP G1500, VSP F1500, VSP, or HUS VM storage systems, selecting this option specifies whether you want to prioritize the relocation of the data in HDT volumes.
Volume Selection	Tier	**Note:** This field is not the same field as the Volume Selection field that is defined when Volume Type is set to Dynamic Tiering. If your selected storage system is configured with storage tiers, you can Select a Tier for volume allocation. If your selected storage system was Any, you can see which storage systems have tiers available using Select a Tier, and make a selection.
	Automatic	Device Manager automatically selects volumes based on the specified volume criteria. For a Basic Volume, Device Manager selects a parity group based on the volume criteria. The selected parity group changes with changes in criteria. If desired, click Select Parity Groups to make your selection. For Dynamic Provisioning volumes, the system selects an HDP pool based on volume criteria. The selected HDP pool changes with the changes in criteria. New volumes may be created or existing volumes may be used. If desired, click Select Pool to make your pool selection. Drive type, speed, chip type, RAID level, and parity group or pool can be adjusted as desired, but the primary purpose of this field relies on the storage system to decide which resources to use for volume allocation.
	Manual	For Basic Volume and Dynamic Provisioning volumes, clicking Manual displays the Selected Volumes dialog box. Click Add Volumes to select basic volumes or Dynamic Provisioning volumes.
Volume Criteria	Drive Type	Indicates the type of drive (FMD, FMC, SSD, SAS, or SATA). The drive type selection can affect volume performance, and the drive type setting can affect the parity group or HDP pool type choices.
	Drive Speed	Indicates the rotational speed (in RPM) of the drive type. This option is not displayed when an FMD, FMC, or SSD drive is selected in Drive Type. The drive speed selection can affect the parity group or HDP pool type choices.
	Chip Type	Indicates the flash memory chip type of the physical drives. However, this option is only displayed if VSP G1000, VSP

Field	Subfield	Description
		G1500, VSP F1500, Virtual Storage Platform or HUS VM is the selected storage system and SSD is the selected drive type. If these two conditions are met, you can select one of three options: Any, SLC, or MLC.
	RAID Level	Changing the RAID level can change the parity group or HDP pool. The RAID levels and parity group configuration can vary depending on the selected storage system. For example, RAID6 may be supported on one storage system, but not on another. In addition, RAID6 is displayed only if RAID6 parity groups have been configured on the selected storage system.
	Parity Groups	Changing the volume criteria can change the system selected parity group. If desired, click Select Parity Groups to make your selection.
	Pool	Click Select Pool to select the HDP pool from which to allocate volumes. Changing the volume criteria can change the system selected HDP pool.
Creating Volume Settings	LDEV ID	Creating Volume Settings fields are only displayed when entered volume criteria requires that a new volume be created. See the previous definition for Volume Capacity field for an example. LDEV ID options for volume creation are displayed if VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, VSP F1500, VSP, or HUS VM are the selected storage system, and volume criteria are such that new volumes need to be created. Select auto or manual to locate or identify an Initial LDEV ID for volume creation.
	Format Type	This field displays formatting methods that are available for the volume to be created, and available for basic volumes on the specified storage system. For example, you might see options for setting either a quick or a basic format. Note that during a quick format, the load might become concentrated on some components, lowering the I/O performance of all hosts that are running in the target storage system.
	Stripe Size	This field displays stripe size options for the volume to be created, as supported by the specified storage system (for example, 64 KB).
Resource Group	-	This field only displays allocating volumes for created users (those not associated with built-in accounts), and allows them to potentially source the volume or volumes from multiple resource groups to which they have access rights, as configured for their account type.
Full Allocation	-	• For VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, or VSP F1500 select Use existing settings to use the full allocation settings of the existing DP volumes. For newly created DP volumes, the full allocation settings are disabled by default. • Select Enable to reserve pages of the specified capacity when allocating volumes. • Select Disable to disable the full allocation settings.

Field	Subfield	Description
Specify a new label	-	Volume labels are searchable, and therefore recommended as a way to find volumes. Select the check box to add a LDEV label. If the target volume is an existing volume with a current label, the new label is applied to the volume.
Initial value	-	The smallest number of sequential numbers. The Initial value is not required, but can be useful for differentiation when allocating multiple volumes. For each volume, the number is in ascending order of LDEV ID.
Reflect these labels to the storage system	-	Reflect these labels to the storage system is checked by default so that naming is consistent between HCS and the storage system itself. If the selected storage system does not support label setting, this item is not displayed.
CLPR for DP volume	-	Select cache logical partition (CLPR) values from the list and assign this value to DP volumes when two conditions are met: the target storage system supports CLPR, and DP volumes are selected. Select either of the following values: Automatically, which assigns the default value if it is a new volume (if it is an existing volume, the CLPR is unchanged), or CLPR, which assigns a CLPR number (0-9...) to the unallocated volume. If you are using a VSP G400, VSP G600, or VSP G800 storage system with a NAS module and the firmware version is earlier than 83-03-25-XX/XX, do not use the CLPR named "NASSystemCLPR". If you use this CLPR, system performance might be severely affected.
Command Device	-	When allocating a volume intended to be a command device for a pair management server, select the Command Device check box and select Enabled. You may also enable or disable the following: • Command Device Security • User Authentication • Device Group Definition
ALUA	-	To enable ALUA, do the following: • Configure the ALUA volume attribute. To enable ALUA volume attribute, select Enable. • Configure Asymmetric Access State of the host group. To set path priorities for asymmetric access to the host group, click Host Group Settings in Detail, and then from the drop-down list, select **Active/Optimized** for primary preferred paths, or **Active/Non-Optimized** for secondary paths, the less preferred paths. In general, the optimized paths are given preference over non-optimized paths.
Use secondary volume of global-active device	-	When primary and secondary storage systems are not discovered by a single HCS management server, select this check box to reserve the secondary volume to prevent access until the global-active device pair can be created using Replication Manager.
>> LUN Path Options	(See the following two	Click LUN Path Options to display the fields and buttons for configuring LUN paths (the storage port to host port mappings that connect volumes to hosts).

Field	Subfield	Description
	fields for details)	
No. of LUN paths per Volume	-	Specify the number of LUN paths to allocate per host. Changing the path count may cause the system to suggest a new path for you automatically. Click Edit LUN Paths to assign or change LUN paths. LUN paths can be displayed in either graphical (Topological Graph) or tabular (Selection Table) formats. In both views, use these links to display WWN nickname information to confirm the target host bus adapter (HBA). • In the graphical view, click on a storage port row to add it to the **LUN Path Editor** panel. Connect the line to the target HBA. When another line is displayed, you can connect it to another HBA or discard it with a click. • In the tabular view, select a storage port, select a host port row, and then click Add to move the mapping to the Selected host Ports panel. When editing the LUN path of a global-active device paired volume, you can specify the settings while referencing the LUN paths of the other volume. To delete default or incorrect mappings, click the connector line in the graphical view, or click Remove in the tabular view. For NAS Platform family, the path redundancy setting is recommended. By default, the number of LUN paths displayed is equal to the number of NAS Platform ports or storage system ports, whichever is lower.
LUN Security Disabled Storage Ports, Selected Storage Ports	-	These related screens only display under LUN Path Options when Allocate without LUN security is selected from the host drop-down. For more information, see the previous Hosts field description. Unsecured ports are listed in LUN Security Disabled Storage Ports. Select the ports on which to make the volume available, and click Add to move the ports to Selected Storage Ports. You can review your plan and click Submit for this allocation task. All hosts that can access the ports on which the volume is on can access the volume.
>> Host Group and LUN Settings	(See the following two fields for details)	Volume allocations using Fibre Channel or FCoE prompt for such items as host group name, host mode, host mode options, and LUN number. To display or hide the following fields, click Host Group and LUN Settings.
Host Groups	Shared by All Hosts	This option specifies the configuring of all hosts in a single host group when volumes are allocated to multiple hosts simultaneously.
	Separate for Each Host	This option specifies the configuring of a separate host group for each host when volumes are allocated to multiple hosts simultaneously.
	Name	Displays host group information for hosts that are in a host group. For hosts that are not in a host group, a host group can be created by entering a name for a host group as needed.

Field	Subfield	Description
	Host Group Settings in Detail	This button displays a list of host groups than can be used during volume allocation, or which were previously created for prior volume allocations.
	Host mode	Select the host mode that supports the host type for which you are allocating volumes.
	Host mode options	Select one or more host mode options for supporting special requirements for specific applications.
LU Number	auto or manual selection buttons	Logical unit (LU) numbers can be assigned automatically or manually for the volumes being allocated. To automatically assign LU numbers to the allocated volumes, select auto and enter a start from number (or use the default). LU numbers are set in ascending order while avoiding existing numbers. To manually assign LU numbers to the allocated volumes, select manual and click Select LU Number to choose a starting LU number.
>> Virtual ID Settings	(See the next field for details)	Virtual ID Settings for global-active device are displayed only when options other than the default virtual storage machine are used on the primary storage. An option for manually specifying the starting virtual LDEV ID to use for allocated volumes. Volumes for which virtual IDs have not been specified are assigned a new virtual LDEV ID when they are allocated to a host that belongs to resource groups used in data migrations that use virtual IDs. Virtual LDEV IDs are assigned automatically by default, with unused IDs assigned to volumes in ascending order. If a user manually specifies an ID, volumes receive the lowest available ID that is equal to or greater than the specified value. For global-active device, accept the default starting virtual LDEV ID for global-active device paired volumes being allocated, or manually select a starting virtual LDEV ID. The displayed value is the minimum value that can be specified as an LDEV ID. The virtual LDEV IDs will provide a single ID for the global-active device paired volumes being accessed by hosts.
Starting virtual LDEV ID Targets	LDKC	Logical disk controller (LDKC) number that forms part of the starting virtual LDEV ID.
	CU	Control unit (CU) number that forms part of the starting virtual LDEV ID.
	DEV	Device (DEV) number that forms part of the starting virtual LDEV ID.
>> Pair management Server Settings	(See the following two fields for details)	These selections are made on the Primary and Secondary tabs.
Pair management server	-	Displayed hosts will be pair management servers configured during global-active device setup.

Field	Subfield	Description
Instance ID	New	Create a new configuration definition file (with a new instance ID). The entered instance ID is validated to prevent duplicates. You can also specify a specific UDP port number for communications. Available UDP port numbers are: • 0 - 899 • 999 - 4093
	Existing	Use an existing configuration definition file. The existing instance IDs are listed.
>> Pair Settings	(See the following 3 fields for details)	For global-active device, pair settings fields are used to finalize pair information.
Quorum Disk	-	The quorum disk number is configured for the primary and secondary storage systems during the initial global-active device setup. You can accept the default value or select a value.
CTG ID	-	When you enable this option you can select the CTG ID to be assigned for the global-active device pair. To add the GAD 3DC delta resync configuration to allocated global-active device pair volumes, CTG IDs need to be assigned.
Copy Group	New	Used to create a new copy group by name for the global-active device pairs being allocated. Valid copy group names must be unique in a configuration definition file (horcmN.conf) and follow these rules: • 1 - 31 characters • A - Z, a - z, 0 - 9, (dash -), (underscore _), (period .), @ • '-' is not permitted as the first character
	Existing	Places global-active device pairs into an existing copy group. Existing is disabled if no copy group exists.
Pair name	Automatic	A pair name consisting of a prefix and start sequence no. is created automatically by the system.
	Manual	A pair name consisting of prefix and start sequence no. is created manually by the user. Valid pair names must be unique in a copy group, and follow these rules: • 1 - 26 characters • A - Z, a - z, 0 - 9, (dash -), (underscore _), (period .), @ • '-' is not permitted as the first character • Start sequence no should be 0 - 99999 • The prefix and start sequence no. should not be blank
>> iSCSI Target and LUN Settings	(See the following two fields for details)	Volume allocations using iSCSI will prompt for items such as host mode, host mode options, name of the iSCSI target, and LU number. To display or hide the following fields, click iSCSI Target and LUN Settings.
iSCSI Targets	Shared by All Hosts	Select this option to indicate that all hosts in the volume allocation reference the same iSCSI targets.
	Separate for each Host	Select this option to indicate that each host in the volume allocation references different iSCSI targets.

Field	Subfield	Description
	Selecting Method	Select Use an existing iSCSI Target to indicate that you want to use existing iSCSI targets. Select Create an iSCSI Target to indicate that you want to create new iSCSI targets.
	Name	Prompts you to create a new iSCSI target, or displays the existing iSCSI target name.
	iSCSI Target Settings in Detail	Displays a list of iSCSI targets that can be used during allocation. This also displays iSCSI targets that were set after the task for allocating volumes was executed.
	Host mode	Select the host mode that supports the host type for which you are allocating volumes. **Note:** When volumes are allocated to multiple hosts whose host modes differ, Mixed is displayed and indicates an error.
	Host mode options	Select one or more host mode options to support special requirements for specific applications.
LU Number	-	This is the same as described in Host Groups.
>> Host Group and LUN Settings for Mid-range Storage	(See the following field for details)	For mid-range storage systems, the Host Group and iSCSI target dialog boxes can include additional fields for enhancing the management of mid-range storage systems. The following fields are specific to mid-range storage systems only. The Host Mode and LU Number fields remain part of the mid-range storage dialog box and are documented in the previous descriptions.
Options	Platform	Selects the host platform type (for example, Windows) to assist in setting host mode.
	Middleware	Selects the installed software on the platform to assist in setting host mode.
	Alternate Path	Indicates the multipath management software used by the host to assist in setting host mode.
	Failover	Indicates the cluster software used by the host to assist in setting host mode.
	Additional Parameters	Click Select to display additional specific usage options to assist in setting host mode.

Related concepts

- About allocating volumes on page 201

Related tasks

- Allocating volumes from general tasks on page 204
- Allocating volumes to selected hosts on page 205
- Allocating volumes to selected file servers on page 206
- Allocating selected volumes to hosts on page 208
- Allocating volumes to clustered hosts on page 209
- Allocating volumes by using a keyword search on page 210
- Allocating volumes by using a criteria search on page 211
- Allocating volumes by using existing volume settings on page 212
- Allocating global-active device pairs on page 350

- Allocating global-active device pairs based on like volumes on page 353
- Allocating global-active device pairs using allocated volumes on page 356
- Configuring pair management servers on page 348

Related references

- Volume allocation methods on page 202
- Notes on performing quick formats on page 151

About clustered-host storage

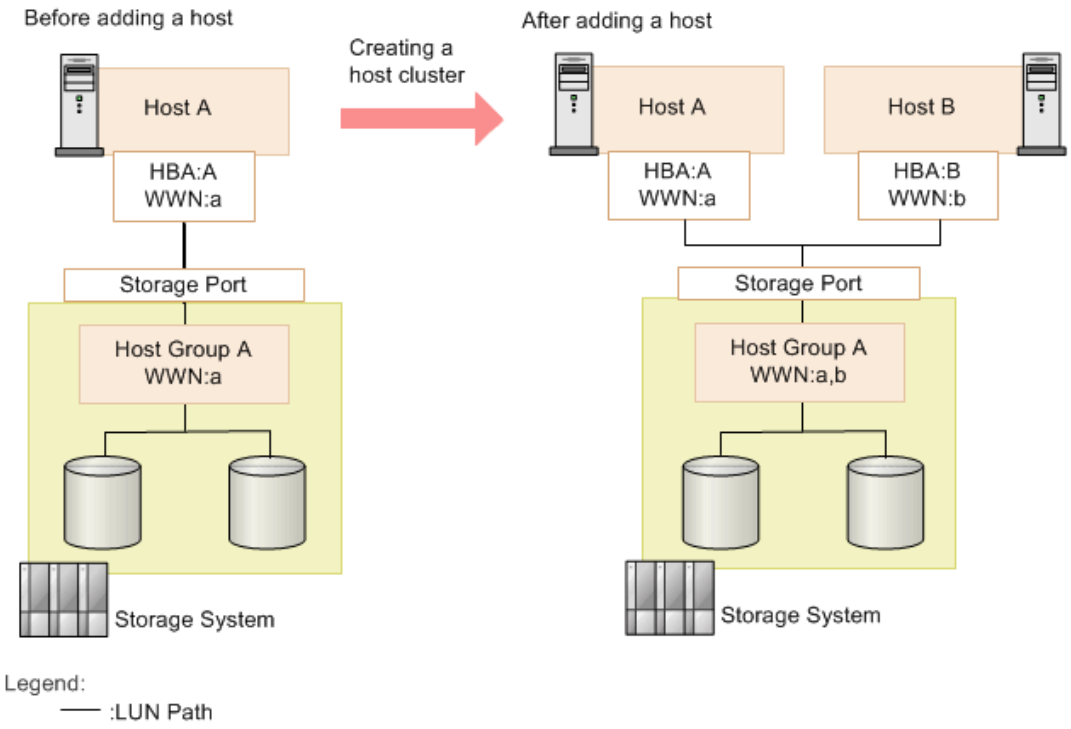
Clustered-host storage is a storage configuration that is created when volumes are allocated to a new host (or file server) that is added to a host group (also known as a host cluster).

When creating clustered-host storage, you add the WWN of the newly added host to the host group to which the WWN of an existing host belongs, and you set LUN paths from the newly added host to the same volumes as those for an existing host.

For example, to better manage and distribute the load on your applications and resources, update the existing host group by creating clustered-host storage using existing volumes by allocating them to a new host in the host group.

Newly allocated volumes represent additional storage resources for a new host. Clustered-host storage supports the reallocation of existing volumes within the host group to meet specific needs.

The following figure illustrates the process of adding a host to create clustered-host storage in a system.



Related concepts

- About allocating volumes on page 201
- About managing WWNs by using nicknames on page 488
- About LUN path management on page 238

Related tasks

- Creating clustered-host storage on page 224
- Allocating volumes to selected hosts on page 205

Creating clustered-host storage

You create clustered-host storage by allocating volumes to a new host within an existing host group (also known as a host cluster).

Creating clustered-host storage involves allocating new or additional volumes to a new host that has been added to an existing host group, setting new LUN paths to the new host, and adding the WWN of the new host to an existing host group.

Before you begin

- Discover (and register) new hosts.
- Allocate volumes to existing hosts.
- Verify that the host connections are Fibre Channel or Fibre Channel over Ethernet (FCoE)

Procedure

1. On the **Resources** tab, click **Hosts**, and select **All Hosts** to list the hosts by OS type in the Application pane.
2. From the list of hosts, select one or more hosts (or from the list of volumes, select one or more volumes allocated to these hosts).
3. Click **More Actions**, and select **Define Clustered-Host Storage**.
4. Select secondary hosts to add to the cluster:
   - To add one host, select the host name in the **Select a host/file server** list.
   - To add multiple hosts, click **Select Hosts/File Servers**, highlight each host name, and click **Add**.
5. From the **Selected Storage System** list, select the storage system to associate with the selected host.

   If two or more storage systems are allocated to the selected host, select only one storage system.
6. In the **Add WWN** list, select the WWN of the new host to add to an available host group.
   View the host groups to which the new host can be added in the **Available Host Groups** list.
7. In the **Available Host Groups** list, select a host group, and click **Add**.
   You can see the volumes that are associated with the selected host group in **A list of affected volumes**, and also verify the added LUN paths for each entry in the **Available LUN Path** list.

   Note: To add more new hosts and associated WWNs, repeat steps 4 through 7 as needed. If you need to modify host information, expand the **Host Group and LUN Settings** to set host group information.
8. Click **Show Plan** and confirm that the information in the plan summary is correct.
   If changes are required, click **Back**.
9. (Optional) Update the task name and provide a description.
10. (Optional) Expand **Schedule** to specify the task schedule.
    You can schedule the task to run immediately or later. The default setting is **Now**. If scheduled for **Now**, select **View task status** to monitor the task after it is submitted.
11. Click **Submit**.
    If the task is scheduled to run immediately, the task begins.
12. You can check the progress and the result of the task on the **Tasks & Alerts** tab.
    Click on the task name to view details of the task.

Result

The new host is added to the designated host group, which created clustered-host storage by virtue of the following:

- LUN paths are created between new host and host group
- WWN of new host is added to host group

Tip: You can also use the Edit LUN Paths dialog box to confirm that the WWN of the new host is successfully added to the host group.

Related concepts

- About clustered-host storage on page 223

Related tasks

- Allocating volumes to selected hosts on page 205
- Editing a WWN nickname on page 490
- Editing LUN paths on page 240

About unallocating volumes

You can unallocate volumes from hosts or file servers.

You can also unallocate global-active device paired volumes from hosts.

Unallocated volumes can be:

- Re-allocated (with existing data) to a host or file server that can recognize the existing data (file system).
- Used for other storage requirements.

Unallocating a volume deletes all LUN paths that connect the volumes to the host or file server.

Unallocating volumes does not delete existing volume data by default. However, during unallocation there is an option to delete the volume (volume data is lost) and return it to unused capacity, or you can delete the unallocated volume later. As a precaution, to retain volume data, back up volumes before unallocating them (for example, for volumes to be re-allocated to a new host).

Related concepts

- About removing hosts and related resources, or vCenters on page 79

Related tasks

- Unallocating volumes from hosts on page 227
- Unallocating volumes from file servers on page 228

Related references

- Unallocate volumes dialog box on page 229

Unallocating volumes from hosts

Unallocated volumes can be re-allocated (with existing data) or can be made available for other storage requirements.

Before you begin

- Identify the name of the target host, and the volumes to unallocate.
- If necessary, backup data on the target volumes.
- Unmount all allocated volumes that you plan to unallocate. An IT administrator might have to perform this task.

Procedure

1. On the **Resources** tab you can unallocate volumes from several locations:
   - From **General Tasks**, select **Unallocate Volumes**.
   - Select a host OS, select one or more target hosts, and click **Unallocate Volumes**.
   - Select a host OS, click a target host name to display volumes, select one or more volumes, and click **Unallocate Volumes**.
   - Search for a host, click the host name and go directly to the volume list, and click **Unallocate Volumes**

   The Unallocate Volumes dialog is launched.
2. Select the host and host volumes to unallocate. Note that if the host and host volumes were selected prior to launching the Unallocate Volumes dialog box, you will go directly to the plan summary mentioned in the next step.
3. Click **Show Plan** and confirm that the information in the plan summary is correct. If changes are required, click **Back**.
4. (Optional) Update the task name and provide a description.
5. (Optional) Expand **Schedule** to specify the task schedule.

   You can schedule the task to run immediately or later. The default setting is **Now**. If the task is scheduled to run immediately, you can select **View task status** to monitor the task after it is submitted.
6. Click **Submit.**
   If the task is scheduled to run immediately, the process begins.
7. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.

Result

Unallocated volumes are added back to the storage system Open-Unallocated volume list.

Related concepts

- About unallocating volumes on page 226

Related tasks

- Shredding volume data on page 155
- Deleting unallocated volumes on page 157

Related references

- Unallocate volumes dialog box on page 229

Unallocating volumes from file servers

Unallocated volumes can be re-allocated (with existing data) or can be made available for other storage requirements.

Before you begin

- Identify the name of the target cluster or file server and the volumes you want to unallocate.
- If necessary, back up data on the target volumes.
- Unmount all allocated volumes that you plan to unallocate. An IT administrator might have to perform this task.

Procedure

1. On the **Resources** tab, select **File Servers**, then select **All File Servers**.
   - From the **Servers/Clusters** list, select the row of the target cluster or file server (only one row can be selected), and then click **Unallocate Volumes**.
   - To unallocate individual volumes, select the target cluster or file server in the tree, and then select the target volumes from the **System Drives** tab for NAS Platform, or the **Volumes** tab for NAS Platform F or Data Ingestor, and then click **Unallocate Volumes**.
2. Specify the appropriate settings for creating a plan.
3. Verify the plan and click **Submit**.
4. In the task list, confirm that the task is completed.



Tip: For NAS Platform family, for reasons such as node deletion, if you want to unallocate volumes from individual file server nodes that are in clusters, click the link of the cluster name and select the **Physical View** tab. From the list, select the row of the target file server (node) and click the **Unallocate Volumes** button.

Related concepts

- About unallocating volumes on page 226

Related tasks

- Shredding volume data on page 155

- Deleting unallocated volumes on page 157

Related references

- Unallocate volumes dialog box on page 229

Unallocate volumes dialog box

Successfully unallocated volumes are placed in the Open-Unallocated folder of the storage system from which they originated. Unallocated volumes can be reallocated to another host, with data intact.

Volumes can be deleted using this dialog box and returned to unused capacity if eligible, or previously unallocated volumes can be reallocated to another host.

When you enter the minimum required information in this dialog box, the Show Plan button activates to allow you to review the plan. Click the Back button to modify the plan to meet your requirements.

The following table describes the dialog box fields, subfields, and field groups. A *field group* is a collection of fields that are related to a specific action or configuration. You can minimize and expand field groups by clicking the double-arrow symbol (>>).

As you enter information in a dialog box, if the information is incorrect, errors that include a description of the problem appear at the top of the box.

Table 14 Unallocate volumes dialog box

Field	Subfield	Description
Host	-	Note: It is recommended that volumes be unmounted from hosts prior to unallocating the volumes. The unallocate volumes dialog box works slightly differently depending on where you launch the dialog box, for example: If you select Unallocate Volumes from General Tasks, you will be prompted to select a host. Selecting the host will display the host volumes. Select one or more volumes and click Show Plan to display the fields and options below. If you select Unallocate Volumes by selecting a host row from the Hosts panel, you will not be prompted to select a host. Host volumes are displayed. Select one or more volumes and click Show Plan to display the fields and options below. If you click the host name to list host volumes, you can identify and select volumes using details such as host group, or volume attribute. Select volumes and click Unallocate Volumes to display the fields and options below. In this case, Show Plan is not displayed because the host volumes were known prior to launching the dialog box.

Field	Subfield	Description
Unallocate global-active device pair simultaneously	-	When a global-active device primary (P-VOL) or secondary (S-VOL) volume is selected to be unallocated and this check box is selected, the global-active device pair will be simultaneously unallocated.
Plan Summary	-	For one or more volumes the volume ID, storage system, volume type, drive type, and host for each volume is displayed. When a global-active device paired volume is selected, Copy Info(P-VOL) and Copy Info(S-VOL) are displayed.
>> Plan Details	Volume Information	There is a Deletable column with a yes or no value to indicate whether the volume can be deleted. For example, volume conditions preventing deletion include: • The volume is in use by another host • The volume is a command device • The volume is a pair volume
	LUN Path Information	For one or more volumes the volume ID, storage system, storage port, port type, host port, WWN/iSCSI name, LU number, host group/iSCSI target, host mode, and 'other options' for each volume is displayed. This provides complete storage and host LUN path information.
	global-active device pair information	When unallocating global-active device pairs, copy pair names, pair management servers, and copy groups appear as information of the pairs that are released at the same time that global-active device pair volumes are unallocated.
	Pair Management Server Information	When unallocating global-active device pairs, information such as the names and instance IDs of the copy groups that are deleted at the same time that pairs are released appears.
	Virtual LDEV ID information	When unallocating volumes such as global-active device S-VOLs, virtual LDEV ID information is displayed.
>> Advanced Options	Host Group/ iSCSI Target Delete	If you select all volumes for a host group/iSCSI target, a host group/iSCSI target deletion option is displayed. Deleting a host group/iSCSI target is only done under very specific circumstances, such as the target server having been replaced with a new server. Do not select this option unless you know the exact status of the target server and the server volumes (data), and are confident this option is appropriate.
	Only unallocate volumes	This default selection unallocates volumes only without deleting volume data or the volume. Unallocated volumes can be re-allocated to another host, for example a newer and faster server.
	Release LUSE volumes	A LUSE volume is created by aggregating (combining) multiple smaller LDEVs into a larger volume that is allocated to a host. If one or more LUSE volumes are selected, this option is activated and can be selected to unallocate LUSE volumes and release the component LDEVs of the LUSE volume.
	Delete volumes	For volumes where Plan Details > Volume Information indicates the volume can be deleted (Deletable=yes), eligible volumes will be deleted and returned to unused

Field	Subfield	Description
		capacity. To delete a selected volume when unallocating volumes: select Delete Volumes, and click Submit.
	Delete virtual ID information assigned to volumes	Select Delete Virtual Information from volumes to delete virtual ID information. This option should only be used by an administrator knowledgeable about the status of migrated volumes with virtual ID information, or the status of global-active device pairs.
	LDEV IDs are moved to the resource pool of the default virtual storage machine	Select LDEV IDs are moved to the resource pool of the default virtual storage machine to release LDEV IDs.

Related concepts

- About unallocating volumes on page 226
- About releasing a LUSE volume on page 159

Related tasks

- Unallocating volumes from hosts on page 227
- Unallocating volumes from file servers on page 228
- Unallocating global-active device pairs on page 360
- Unallocating individual global-active device volumes on page 362

Configuring Fibre Channel and iSCSI ports

Enable LUN security and configure port settings, such as data transfer speeds, port addresses, and fabric switch settings.

Enabling LUN security on a port

To protect mission-critical data in your storage system from illegal access, secure the logical volumes in the storage system. Use LUN Manager to secure LUs from illegal access by enabling LUN security on ports.

By default, LUN security is disabled on each port. When registering hosts in multiple host groups, you must enable LUN security (set the switch to enabled). When you change LUN security from disabled to enabled, you must specify the WWN of the host bus adapter.

Caution: It is best to enable LUN security on each port when configuring your storage system. Although you can enable LUN security on a port when host I/O is in progress, I/O is rejected with a security guard after enabling.

Procedure

1. On the **Resources** tab, click **Storage Systems**, and then expand **All Storage Systems** and the target storage system.
2. Choose one of the following options.
   - For VSP G series or VSP F series:
     Click **Ports/Host Groups/iSCSI Targets**.
   - For other available storage systems:
     From the **Actions** list in the application pane, select **Element Manager**. Refer to the documentation for the native management tool for your storage system.
3. In the **Ports/Host Groups/iSCSI Targets** window, click the **Ports** tab.
4. Select the desired port, and then click **Edit Ports**.
5. Select the **Port Security** check box, and then select **Enable**.
6. Click **Finish**. A message appears, confirming whether to switch the LUN security. Clicking **OK** opens the **Confirm** window.
7. In the **Confirm** window, confirm the settings, in **Task Name** type a unique name for this task or accept the default, and then click **Apply**.
   If **Go to tasks window for status** is checked, the **Tasks** window opens.

Related tasks

- Disabling LUN security on a port on page 232
- Setting the data transfer speed on a Fibre Channel port on page 235
- Setting the Fibre Channel port address on page 237
- Setting the fabric switch on page 237

Disabling LUN security on a port

Caution: Do not disable LUN security on a port when host I/O is in progress.

Procedure

1. On the **Resources** tab, click **Storage Systems**, and then expand **All Storage Systems** and the target storage system.
2. Choose one of the following options.
   - For VSP G series or VSP F series:
     Click **Ports/Host Groups/iSCSI Targets**.
   - For other available storage systems:
     From the **Actions** list in the application pane, select **Element Manager**. Refer to the documentation for the native management tool for your storage system.
3. In the **Ports/Host Groups/iSCSI Targets** window, click the **Ports** tab.
4. Select the desired port, and then click **Edit Ports**.

5. Select the **Port Security** check box, and then select **Disable**.
6. Click **Finish**. If disabling LUN security, a message appears, indicating that only host group 0 (the group whose number is 00) is to be enabled. Clicking **OK** opens the **Confirm** window.
7. In the **Confirm** window, confirm the settings, in **Task Name** type a unique name for this task or accept the default, and then click **Apply**.

   If **Go to tasks window for status** is checked, the **Tasks** window opens.

Related tasks

- Enabling LUN security on a port on page 231

Editing iSCSI port settings

Caution: It is best to enable LUN security on each port when configuring your storage system. Although you can enable LUN security on a port when host I/O is in progress, I/O is rejected with a security guard after enabling.

Procedure

1. On the **Resources** tab, click **Storage Systems**, and then expand **All Storage Systems** and the target storage system.
2. Choose one of the following options.
   - For VSP G series or VSP F series:
     Click **Ports/Host Groups/iSCSI Targets**.
   - For other available storage systems:
     From the **Actions** list in the application pane, select **Element Manager**. Refer to the documentation for the native management tool for your storage system.
3. Select the **Ports** tab.
4. Select the desired port.
5. Display the Edit Ports window by performing one of the following:
   - Click **Edit Ports**.
   - From the **Actions** menu, select **Ports/Host Groups/iSCSI Targets**, then **Edit Ports**.
6. Select the check box to change the option and specify values. Items that can be changed are as follows:
   - **IPv4 Settings**: Specifies **IP Address**, **Subnet Mask**, or **Default Gateway**.
   - **IPv6 Mode**: Specifies enable or disable of this mode.
   - **IPv6 Settings**: Specifies **Link Local Address**, **Global Address**, **Global Address 2**, or **Default Gateway** if **IPv6 Mode** is selected to **Enable**.
   - **Port Security**: Specifies enable or disable.

- **Port Speed**: Specifies the data transfer speed.
- **TCP Port Number**: Specifies the TCP port number.
- **Selective ACK**: Specifies enable or disable.
- **Delayed ACK**: Specifies enable or disable.
- **Maximum Window Size**: Specifies the size of the maximum window.
- **Ethernet MTU Size**: Specifies the MTU size.
- **Keep Alive Timer**: Specifies the keep alive timer.
- **VLAN Tagging Mode**: Specifies enable or disable.
- **iSNS Server**: Specifies enable or disable. If this option is selected to **Enable**, specify **IP Address** or **TCP Port Number**.
- **CHAP User Name**: Specifies the CHAP user name.
- **Secret and Re-enter Secret**: Specifies the secret which is used for host authentication.

7. Click **Finish**. A message appears, confirming whether to switch the LUN security. Clicking **OK** opens the **Confirm** window.
8. In the **Confirm** window, confirm the settings, in **Task Name** type a unique name for this task or accept the default, then click **Apply**. If **Go to tasks window for status** is checked, the **Tasks** window opens.

Setting the T10 PI mode on a port

Before you begin

- The Storage Administrator (Provisioning) role is required to perform this task.
- Access to the CHB(FC) Fibre Channel board port is required.
- The port speed must be 16 Gbps or 32 Gbps.

Caution: If you change the T10 PI mode of one port, the T10 PI mode of the other port paired with changed port also needs to be changed. You must verify the mode of each port in the pair before changing the T10 PI mode. Make sure the ports in each pair are the same in the resource group. The following shows pairs of port IDs.

If you change the setting on one of the ports in the pair, the setting on the other port in the pair will also be changed:

- Port IDs 1*x* and 3*x* (where *x* is a letter from A to R). For example, 1A and 3A can be paired with each other.
- Port IDs 2*x* and 4*x* (where *x* is a letter from A to R). For example, 2B and 4B can be paired with each other.
- If the T10 PI mode of the CHB port of the 4HF32R package, port IDs 1*x*, 3*x*, 5*x*, and 7*x* (where *x* is a letter from A to R) are collectively changed.

Note: If the T10 PI mode is enabled between the path of the target port and LDEV, you cannot disable the T10 PI mode of the port.

Procedure

1. On the **Resources** tab, click **Storage Systems**, and then expand **All Storage Systems** and the target storage system.
2. Choose one of the following options.
   - For VSP G series or VSP F series:
     Click **Ports/Host Groups/iSCSI Targets**.
   - For other available storage systems:
     From the **Actions** list in the application pane, select **Element Manager**. Refer to the documentation for the native management tool for your storage system.
3. Select the **Ports** tab.
4. Select the desired port.
   To collectively change the T10 PI mode of ports, do not intermix ports with enabled or disabled T10 PI modes.
5. Click **Edit Ports**.
6. Click **OK** on the message window.
7. Select **Enable** or **Disable** on the **Edit T10 PI Mode** window.
8. Click **Finish**. The **Confirm** window appears.
9. In the **Confirm** window, confirm the settings. In the **Task Name**, type a unique name for this task or accept the default, then click **Apply**. If **Go to tasks window for status** is checked, the **Tasks** window opens.

Setting the data transfer speed on a Fibre Channel port

As system operation continues, you might notice that a large amount of data is transferred at some ports, but a small amount of data is transferred at other ports. You can optimize system performance on a Fibre Channel port by setting a faster data transfer speed on ports where a large amount of data is transferred, and setting a slower data transfer speed on ports where a smaller amount of data is transferred.



Note: For VSP G1000, G1500, or VSP F1500: In Fibre Channel over Ethernet (FCoE) networks, the port speed is fixed at 10 Gbps and cannot be changed.

Procedure

1. On the **Resources** tab, click **Storage Systems**, and then expand **All Storage Systems** and the target storage system.
2. Choose one of the following options.
   - For VSP G series or VSP F series:
     Click **Ports/Host Groups/iSCSI Targets**.
   - For other available storage systems:
     From the **Actions** list in the application pane, select **Element Manager**. Refer to the documentation for the native management tool for your storage system.

3. In the **Ports/Host Groups/iSCSI Targets** window, click the **Ports** tab.
4. Select the desired port, and then click **Edit Ports**.
5. In the **Edit Ports** window, select the **Port Speed** check box, and then select the desired port speed.

   Select the speed of the Fibre Channel port in the unit of Gbps (Gigabit per second). If **Auto** is selected, the storage system automatically sets the speed to 2, 4, 8, or 16 Gbps.

Caution: Observe the following cautions when setting speed on a Fibre Channel port:

- If the host bus adapters (HBAs) and switches support 2 Gbps, use the fixed speed of 2 Gbps for the channel adapter for Fibre Channel (CHF) or channel board for Fibre Channel (CHB(FC)) or channel board for Fibre Channel (CHB(FC)) port speed. If they support 4, 8, or 16 Gbps, use 4, 8, or 16 Gbps for the CHF or CHB(FC) port speed, respectively.
- If the Auto Negotiation setting is required, some links might not be up when the server is restarted. Check the channel lamp. If it is flashing, disconnect the cable, and then reconnect it to recover from the link-down state.
- If the CHF or CHB(FC) port speed is set to **Auto**, some equipment might not be able to transfer data at the maximum speed.
- When you start a storage system, HBA, or switch, check the host speed appearing in the Port list. If the transfer speed is different from the maximum speed, select the maximum speed from the list on the right, or disconnect, and then reconnect the cable.
- The available port speed which is specified in **Port Speed** is limited due to the combination of the type of the Fibre Channel port and the topology which is specified in **Connection Type**.

---

6. Click **Finish**.
   The **Confirm** window appears.
7. In the **Task Name** text box, type a unique name for the task or accept the default.

   You can enter up to 32 ASCII characters and symbols, with the exception of: \ / : , ; * ? " < > |. The value <*date*>-<*window name*> is entered by default.
8. Click **Apply**.

   If the **Go to tasks window for status** check box is selected, the **Tasks** window appears.

Related tasks

- Setting the Fibre Channel port address on page 237

- Enabling LUN security on a port on page 231

Setting the Fibre Channel port address

When configuring your storage system, set addresses for Fibre Channel ports. When addressing Fibre Channel ports, use AL-PA (arbitrated-loop physical address) or loop IDs as the addresses.

Note: For VSP G1000, G1500, or VSP F1500: In FCoE networks, you do not need to set the address of a Fibre Channel port.

Procedure

1. On the **Resources** tab, click **Storage Systems**, and then expand **All Storage Systems** and the target storage system.
2. Choose one of the following options.
   - For VSP G series or VSP F series:
     Click **Ports/Host Groups/iSCSI Targets**.
   - For other available storage systems:
     From the **Actions** list in the application pane, select **Element Manager**. Refer to the documentation for the native management tool for your storage system.
3. In the **Ports/Host Groups/iSCSI Targets** window, click the **Ports** tab.
4. Select the desired port, and then click **Edit Ports**.
5. In the **Edit Ports** window, select the **Address (Loop ID)** check box, and then select the address.
6. Click **Finish**.
   The **Confirm** window appears.
7. In the **Task Name** text box, type a unique name for the task or accept the default.
   You can enter up to 32 ASCII characters and symbols, with the exception of: \ / : , ; * ? " < > |. The value "date-window name" is entered by default.
8. Click **Apply**.
   If the **Go to tasks window for status** check box is selected, the **Tasks** window appears.

Related tasks

- Setting the data transfer speed on a Fibre Channel port on page 235
- Enabling LUN security on a port on page 231

Setting the fabric switch

When you configure your storage system, specify whether the hosts and the storage system are connected via a fabric switch.

Note: For VSP G1000, G1500, or VSP F1500: In FCoE networks, Fabric is fixed to ON. Therefore, you do not need to set Fabric.

Procedure

1. On the **Resources** tab, click **Storage Systems**, and then expand **All Storage Systems** and the target storage system.
2. Choose one of the following options.
   - For VSP G series or VSP F series:
     Click **Ports/Host Groups/iSCSI Targets**.
   - For other available storage systems:
     From the **Actions** list in the application pane, select **Element Manager**. Refer to the documentation for the native management tool for your storage system.
3. In the **Ports/Host Groups/ISCSI Targets** window, click the **Ports** tab.
4. Select the desired port, and then click **Edit Ports**.
5. Select a check box of **Fabric**, and select **ON** if you set the fabric switch. If you do not set the fabric switch, select **OFF**.
6. Click **Finish**.
   The **Confirm** window appears.
7. In the **Task Name** text box, type a unique name for the task or accept the default.
   You can enter up to 32 ASCII characters and symbols, with the exception of: \ / : , ; * ? " < > |. The value "date-window name" is entered by default.
8. Click **Apply**.
   If the **Go to tasks window for status** check box is selected, the **Tasks** window appears.

Related tasks

- Enabling LUN security on a port on page 231

Managing LUN paths

This module describes path management; how to specify LUN paths, edit host mode and host mode options, and how to specify LUN paths after replacing or exchanging a host bus adapter.

About LUN path management

When you allocate a volume to a host, Hitachi Command Suite allows you to assign or edit LUN paths between one or more volumes and one or more hosts.

LUN paths provide volume access for the host by pairing storage ports and host ports. For example, one or more storage ports can be mapped to one or more host ports or iSCSI targets.

After volumes have been allocated to a host, you can edit existing LUN path information to add new paths or delete existing paths from a list of allocated volumes in the storage systems tree, hosts tree, logical groups tree, or file servers tree, and from a list of volumes returned from a search. However, you cannot edit the LUN paths between NAS modules and User LUs.

To set a new LUN path between a specific volume and a host or file server, perform the allocate volume operation, and to delete all the LUN path settings between a specific volume and a host or file server, perform the unallocate volume operation. In contrast, in LUN path editing, you can change the number of LUN paths and the connection-destination ports according to operations. In particular, edit LUN paths to achieve the following:

- Improvement of I/O performance
  If the usage frequency of an application increases, you can add LUN paths to increase the data transmission speed and improve the I/O performance. If the usage frequency decreases, you can delete the LUN paths.
- Enhancements to system redundancy
  To prepare for an error in storage systems or host ports, you can enhance the system redundancy by configuring multiple LUN paths, with each using a different port.
- Response to failures
  When a port is disabled because of an error, you can configure LUN paths that temporarily use alternate ports to continue system operation.

Because allocated volumes belong to a specific host group or iSCSI target, the same target host port is set for all volumes in any given host group or iSCSI target. Therefore, when adding or deleting a host port for one or more volumes, you must select all of the volumes that belong to that host group or iSCSI target to maintain consistent LUN path assignments for all volumes.

When a FC or FCoE connection is used, you can change host modes and host mode options depending on the situation, for example when an application is added to the host and the operating system is upgraded.

You can add or exchange HBAs to improve performance and throughput requirements. To edit LUN paths when replacing a failed HBA or performing a planned HBA replacement, use any of the following options, depending on your task purpose:

- Add HBA
- Exchange HBA
- Remove HBA

You can use the LUN paths that were set for the old HBA for the new HBA, and delete LUN paths set for an HBA that is no longer necessary. You can also edit LUN paths for multiple HBAs collectively.

Related tasks

- Editing LUN paths on page 240
- Editing the host mode and host mode options on page 242
- Editing LUN paths when exchanging a failed HBA on page 243
- Editing LUN paths when adding or exchanging an HBA on page 244
- Removing LUN paths after adding an HBA on page 245

Editing LUN paths

You can manage the LUN paths between storage systems and hosts by adding or deleting them as needed.

You manage LUN paths on your storage system by controlling the connections between the storage systems and hosts. This allows you to better adapt to changing storage system and network conditions.

As conditions change, you can create new paths or delete selected paths on a LUN path basis that are established between multiple host bus adapters (HBA) of a host within a host group. For example, some HBAs may become unnecessary when the related applications are unneeded or infrequently used, and you can delete selected LUN paths for such cases.

Before you begin

Allocate volumes to the existing hosts.

Procedure

1. On the **Resources** tab, expand the tree for storage systems, hosts, file servers, or logical groups to display volumes.
2. Select one or more volumes for which you want to edit LUN paths. If selecting multiple volumes, they must belong to the same configuration host group.
3. In the selected volume, click **More Actions**, and select **Edit LUN Paths**.
4. In the **Edit LUN Paths** dialog box, use the topographical graph or selection table view to map storage ports to host ports. In both views, you can use links to view WWN nickname information to confirm the target HBA.

   Tip: When editing the LUN path of a global-active device paired volume, you can specify the settings while referencing the LUN paths of the other volume.

   a. In **Topological Graph** (or graph view), click on a storage port row to add it to the **LUN Path Editor** panel. Connect the line to the target

HBA. Another line is displayed, which you can connect to another HBA or discard with a click.

b. In **Selection Table** (or table view), first select a storage port, then select a host port row, and click **Add** to move the mapping to **Selected host Ports** list.

If you add an incorrect mapping and want to delete it, click the connector line in graph view, or use the remove button in table view, or click **Cancel** to close the dialog box and start over.

5. To delete an existing path, which is indicated by a green line in graph view, or by the 'In Use' state in the Selected host ports list in **Selection Table**, do the following:

Tip: To delete all the LUN paths between a specific volume and a host, file server so that there are no paths left, delete LUN path settings by unallocating the volumes.

a. In graph view, click the green line. The line will now be gray and thin which indicates that the LUN path is removed.

b. In table view, select the mapping row and click **Remove** to change the state from 'In Use' to 'Remove.'

6. Change any other required settings.
7. Click **Show Plan** and confirm that the information in the plan summary is correct.
   Optionally, update the task name and provide a description.
8. Expand **Schedule** to specify the task schedule.
   The task can be run immediately or scheduled for later. The default setting is **Now**.
9. Click **Submit**.
   If you selected **Now**, the editing LUN path process begins
10. You can check the progress and result of the editing LUN paths task on the **Tasks & Alerts** tab.
    Verify the results for each task by viewing the details of the task.

Result

The LUN path settings you edited are displayed correctly.

Related concepts

- About LUN path management on page 238

Related tasks

- Allocating global-active device pairs on page 350
- Allocating global-active device pairs based on like volumes on page 353
- Allocating global-active device pairs using allocated volumes on page 356

Editing the host mode and host mode options

Manage your host group information by editing the host mode and host mode options after the volumes have been allocated.

You can edit host group information (host group name, host mode, and host mode options) for an existing host group by editing its LUN paths when a host, to which volumes are allocated, has been added to a host cluster. You can verify host mode option IDs when host group information is being edited or when new volumes are allocated.



Tip: In addition to when you are editing LUN paths, you can also edit host group information when allocating volumes, allocating like volumes, and defining clustered-host storage.

Before you begin

- Allocate volumes to hosts.
- Verify that the host connections are Fibre Channel or Fibre Channel over Ethernet (FCoE).

Procedure

1. On the **Resources** tab, choose one of the following options to expand the tree for storage systems, hosts, file servers, or logical groups to display volumes.
   - **Storage Systems > All Storage Systems**
   - **Hosts > All Hosts**
   - **File Servers > All File Servers**
   - **Logical Groups > Public Logical Groups** or **Private Logical Groups**
2. Select a volume that belongs to the target host group (if selecting multiple volumes, they must belong to the same configuration host group).
3. Click **More Actions**, and select **Edit LUN Paths**.
4. In the **Edit LUN Paths** dialog box, click **Host Group and LUN Settings**, and modify the host group name, host mode, or host mode options as needed.
5. Click **Show Plan** and confirm that the information in the plan summary is correct.
   If changes are required, click **Back**.
6. (Optional) Update the task name and provide a description.
7. (Optional) Expand **Schedule** to specify the task schedule.
   You can schedule the task to run immediately or later. The default setting is **Now**. If scheduled for **Now**, select **View task status** to monitor the task after it is submitted.
8. Click **Submit**.
   If the task is scheduled to run immediately, the task begins.

9. (Optional) You can check the progress and the result of the task on the **Tasks & Alerts** tab.
   Click on the task name to view details of the task.
10. (Optional) In the **Host Group and LUN Settings** dialog box, click **Edit LUN Paths** to verify the host mode or host mode options.

Related concepts

- About changing host settings and information on page 73

Related tasks

- Editing LUN paths on page 240
- Creating clustered-host storage on page 224

Editing LUN paths when exchanging a failed HBA

Exchange a failed HBA with a new HBA and restore the LUN path settings to the new HBA.

Before you begin

- Identify the new WWN for the HBA that is being added
- Identify the WWN from which to model paths
- Verify that the new HBA is physically connected.

Procedure

1. On the **Resources** tab, select **Hosts**.
2. After selecting the target operating system, select the target host row, and click **More Actions > Exchange HBAs**.
3. Enter the **New WWN** or select a WWN from the table.
4. Enter the **WWN from which to model paths** or select a WWN from the list.
   The selected WWN will be removed from the host.
5. Click **Add.**
6. In the **WWN Pairs** list, verify that the listed HBA WWN combination before and after the replacement are paired correctly.



Tip:

- If the WWN information is updated when the host is refreshed, the target WWN might not be displayed in the list. In this case, you need to manually enter the WWN of the failed HBA.
- To edit a WWN nickname from list of **WWN Pairs**, click **Edit WWN Nicknames**.

7. Click **Show Plan** and confirm that the information in the plan summary is correct. If changes are required, click **Back**.
8. (Optional) Update the task name and provide a description.

9. (Optional) Expand **Schedule** to specify the task schedule.
   You can schedule the task to run immediately or later. The default setting is **Now**. If the task is scheduled to run immediately, you can select **View task status** to monitor the task after it is submitted.
10. Click **Submit.**
    If the task is scheduled to run immediately, the process begins.
11. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.

Result

When the task completes, the LUN path settings are restored to the new HBA and the original WWN is removed from the host.

Related concepts

- About LUN path management on page 238
- About managing WWNs by using nicknames on page 488

Related tasks

- Editing a WWN nickname on page 490

Editing LUN paths when adding or exchanging an HBA

You can add or exchange HBAs to improve performance and throughput requirements. When adding an HBA, specify the WWN of the new HBA and then select a WWN of an existing HBA from which to model paths.

Before you begin

- Identify the new WWN for the HBA that is being added
- Identify the WWN from which to model paths
- Verify that the new HBA is physically connected.

Procedure

1. On the **Resources** tab, select **Hosts**.
2. After selecting the target operating system, select the target host row, and click **More Actions > Add HBAs**.
3. Enter the **New WWN** or select a WWN from the list.
4. Enter the **WWN from which to model paths** or select a WWN from the list.
5. Click **Add.**
6. In the WWN Pairs list, verify that the listed HBA WWN combination before and after the replacement are paired correctly.

Tip:

- If the WWN information is updated when the host is refreshed, the target WWN might not be displayed in the list. In this case,

you need to manually enter the WWN of the HBA you are adding.
- To edit a WWN nickname from list of **WWN Pairs**, click **Edit WWN Nicknames**.

7. Click **Show Plan** and confirm that the information in the plan summary is correct. If changes are required, click **Back**.
8. (Optional) Update the task name and provide a description.
9. (Optional) Expand **Schedule** to specify the task schedule.
   You can schedule the task to run immediately or later. The default setting is **Now**. If the task is scheduled to run immediately, you can select **View task status** to monitor the task after it is submitted.
10. Click **Submit.**
    If the task is scheduled to run immediately, the process begins.
11. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.

Result

When the task completes, the new WWN is added and related LUN path settings are restored to the host.

Next steps

If you are performing a planned HBA replacement, remove any unnecessary WWNs and LUN path settings.

Related concepts

- About LUN path management on page 238

Related tasks

- Removing LUN paths after adding an HBA on page 245
- Editing a WWN nickname on page 490

Removing LUN paths after adding an HBA

Remove a WWN from the host and also delete the related LUN paths.

Before you begin

- Identify the WWN of the HBA you are removing

Procedure

1. On the **Resources** tab, select **Hosts**.
2. After selecting the target operating system, select the target host row, and click **More Actions > Remove HBAs**.

3. Enter the **WWN** to be removed from the host or select a WWN from the list.
4. (Optional) Select the check box **Delete Host Group** to delete the selected host group. By default, the check box is clear.
5. Click **Show Plan** and confirm that the information in the plan summary is correct. If changes are required, click **Back**.
6. (Optional) Update the task name and provide a description.
7. (Optional) Expand **Schedule** to specify the task schedule.

   You can schedule the task to run immediately or later. The default setting is **Now**. If the task is scheduled to run immediately, you can select **View task status** to monitor the task after it is submitted.
8. Click **Submit.**
   If the task is scheduled to run immediately, the process begins.
9. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.

Result

When the task completes, the WWN and related LUN path settings are removed from the host.

Related concepts

- About LUN path management on page 238

Related tasks

- Editing LUN paths when adding or exchanging an HBA on page 244

Managing NAS Platform file systems and shares

This module describes how to create file systems and shares for NAS Platform file servers and NAS modules.

About managing NAS Platform file systems and shares

The NAS Platform family provides high performance file server capabilities using storage system volumes.

The NAS Platform family, which is managed using Systems Management Unit (SMU)/NAS Manager, includes the following models:

- NAS Platform 3080/3090
- NAS Platform 3100/3200
- NAS Platform 4060/4080/4100

Note: Model references are used in this section only if necessary. For example, to describe features that are supported only on specific models.

When NAS Platform is registered to Device Manager, you can use the volume of the storage system to create storage pools and file systems, and provide file shares by using the Device Manager GUI. On the Device Manager dashboard, you can monitor the usage rate of storage pools and file systems, as well as check NAS Platform alerts.

The operations that can be performed by using the Device Manager GUI differ depending on the versions of NAS Platform or SMU/NAS Manager the system is linked to. Operations that are not supported by Device Manager can be performed by launching SMU/NAS Manager.

Managing storage pools

Create a storage pool using volumes (system drives) allocated to the file server, and use the storage pool for file systems or cache.

The configuration of the storage pool is affected by the file system performance and reliability, capacity efficiency, and expandability. By using the Device Manager GUI, you can apply best practices and easily configure storage pools.

Storage pools can be configured not only from basic volumes, but also from DP volumes. However, note that both a basic volume and DP volume cannot coexist within the same storage pool. If a DP volume is used as a system drive, the I/O of the parity group is equalized due to the rebalancing functionality of Dynamic Provisioning. As a result, the storage pool will sufficiently perform.

You can check the usage of the storage pool from the **Top 10 File Server Storage Pool** report in File Server view or on the dashboard. Expand the storage pool when the used capacity of the file system increases and if the storage pool capacity becomes insufficient.

If you delete a storage pool that is no longer necessary by using Device Manager, volumes that are related to that storage pool are also deleted.

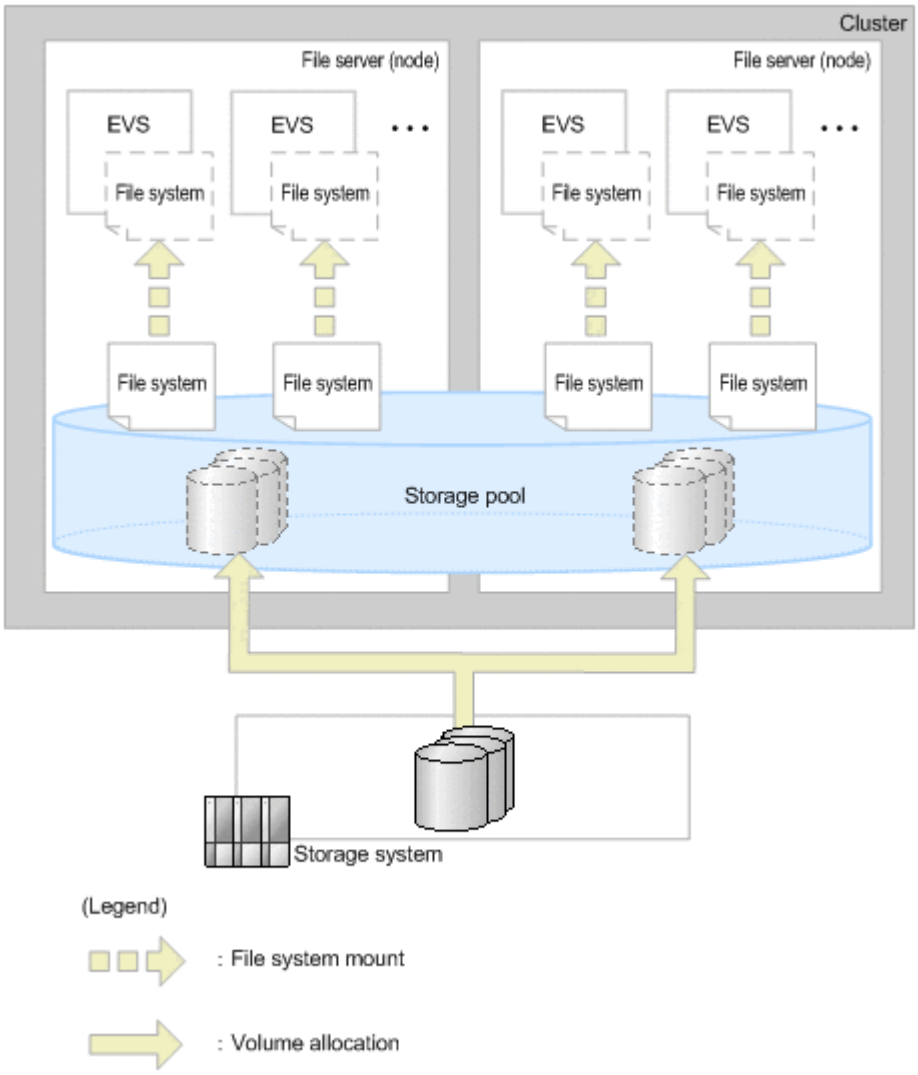
The following table lists the operations that can be performed for storage pools by using Device Manager, and the required version of the file server.

Function	Required version
Create/expand a storage pool in a DP pool	NAS Platform family v12.1 or later
Create/expand a storage pool in a basic volume Delete a storage pool Change a storage pool name	NAS Platform family v11.3 or later
Create/expand/delete a storage pool (launch SMU/NAS Manager)	NAS Platform family v8.1.2351 and SMU/NAS Manager v10.1.3070 or later

Function	Required version
Change the settings of a storage pool (launch SMU/NAS Manager)	

Managing file systems

You can create a file system in a storage pool that consists of a collection of volumes (system drives) allocated to the file server, and mount the file system on Enterprise Virtual Server (EVS).



You can use **Top 10 File Systems** from the dashboard to check for file systems with a high capacity usage. If a file system is at risk of not having enough free capacity, you can add capacity:

- File systems whose allocation type is On Demand Allocation are allocated an appropriate amount of capacity as needed. File systems expand automatically until reaching a predefined maximum size. If the file system reaches this maximum size, you can manually set a higher maximum size.

- File systems whose allocation type is Full Allocation should be expanded manually. When expanding a file system, specifying a capacity immediately allocates that capacity to the file system.

File systems can be mounted and unmounted from Device Manager when you want to copy a file system or perform maintenance operations, such as changing EVS settings.

If you delete file systems that are no longer necessary by using Device Manager, all file shares that are related to the target file system will also be deleted. If you delete a storage pool, all the volumes that are related to the target storage pool will also be deleted.

The following table lists the operations that can be performed for the file system by using Device Manager, and the required version of the file server.

Function	Required version
Expand a file system (Full Allocation type)	NAS Platform family v12.0 or later
Expand a file system (On Demand Allocation type) Mount/unmount a file system Delete a file system	NAS Platform family v11.1 or later
Create a file system	NAS Platform family v10.2.3071 or later
Create/expand/delete a file system (launch SMU/NAS Manager) Change the settings of a file system (launch SMU/NAS Manager)	NAS Platform family v8.1.2351 and SMU/NAS Manager v10.1.3070 or later

Managing File Shares

You can add file shares (CIFS Share or NFS Export) to file systems with adequate free capacity for your purpose. If a file system does not have adequate free capacity, it may be possible to expand the file system.

You can create file shares at the same time as when you create the file system.

A license related to file sharing functionality for NAS Platform is required.

The following table lists the operations that can be performed for the file shares by using Device Manager, and the required version of the file server.

Function	Required version
Create/delete a file share Change the settings of a file share	NAS Platform family v11.1 or later

Function	Required version
Add/delete a file share (launch SMU/NAS Manager) Change the settings of a file share (launch SMU/NAS Manager)	NAS Platform family v8.1.2351 and SMU/NAS Manager v10.1.3070 or later

Related concepts

- Allocating storage on page 141
- About registering and removing file servers on page 81

Related tasks

- Creating a NAS Platform file system on page 264
- Adding NAS Platform file shares on page 270
- Expanding a NAS Platform file system on page 266
- Changing NAS Platform file share settings on page 271
- Creating a NAS Platform storage pool (v11.3 or later) on page 257
- Expanding a NAS Platform storage pool (v11.3 or later) on page 259
- Creating a NAS Platform storage pool (earlier than v11.3) on page 261
- Expanding a NAS Platform storage pool (earlier than v11.3) on page 262
- Changing NAS Platform storage pool settings on page 263

Related references

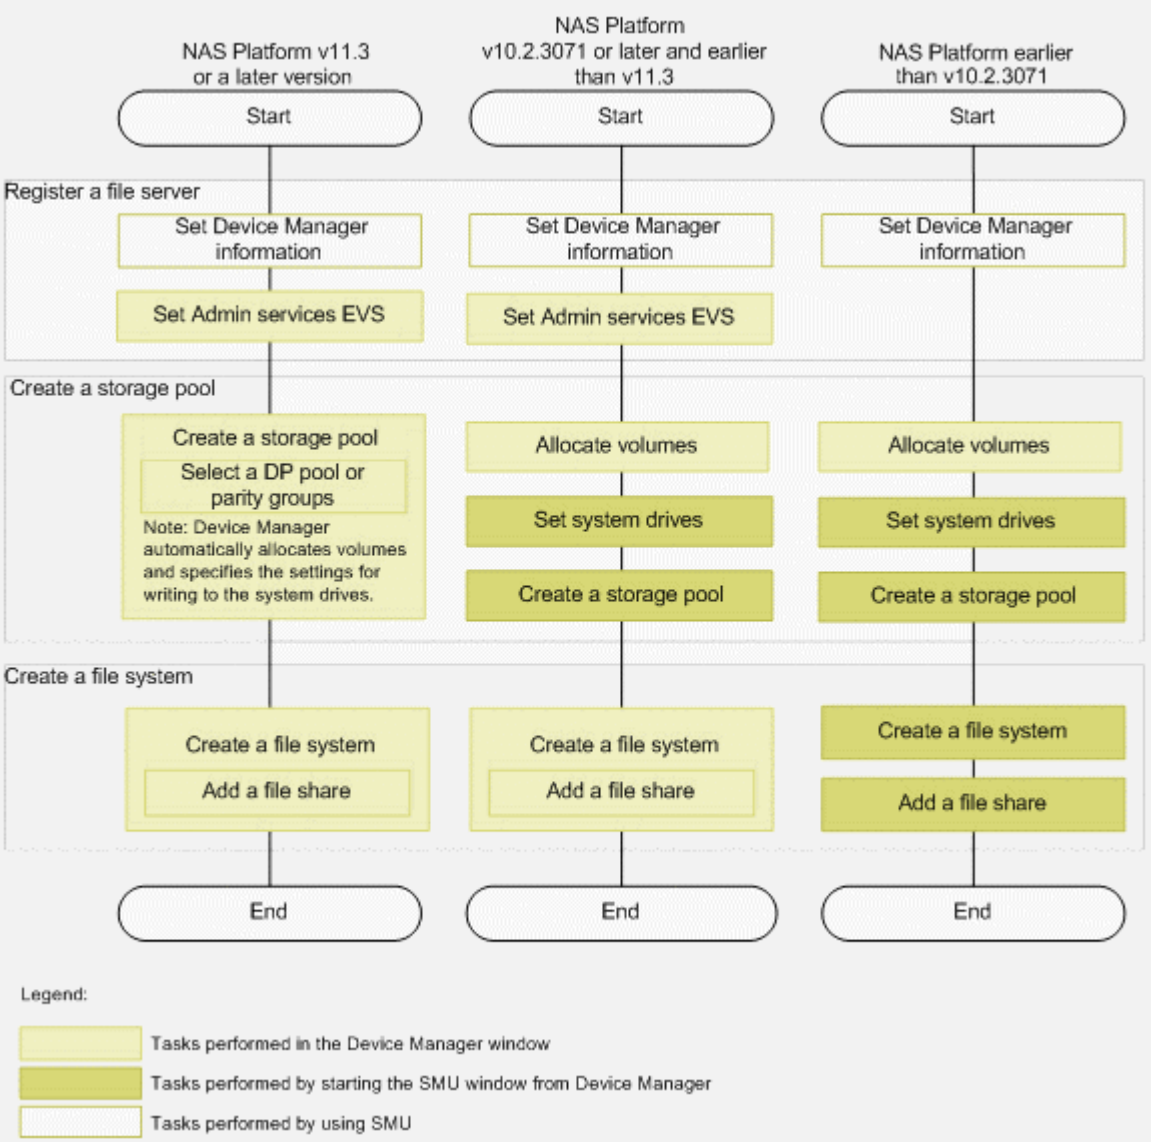
- Workflow for providing NAS Platform file shares on page 250
- File share configuration examples on page 279
- Conditions for parity groups that are used in storage pools on page 253
- Conditions for DP pools that are used in storage pools on page 253

Workflow for providing NAS Platform file shares

Use Device Manager to create files systems and files shares (CIFS or NFS Export) when a NAS Platform file server or NAS module is managed using System Management Unit (SMU)/NAS Manager v10.1.3070 or later.

Tip: When a file server is managed by an SMU version earlier than v10.1.3070, first allocate volumes to the file server by using Device Manager, and then create file systems by using SMU.

The workflow for registering a file server to providing file systems and shares is indicated in the following figure:



- Register a NAS Platform
  - Set Device Manager information
    By setting Device Manager information in SMU/NAS Manager, NAS Platform (file server/NAS module) information that is managed in SMU/NAS Manager is sent to Device Manager.
  - Set Admin services EVS
    For NAS Platform v10.2.3071 or later, using Device Manager, set Admin services EVS information (IP address and user account for the Server Control (SSC) of the file server) for each cluster. For NAS modules that are installed with VSP G400, G600, G800, and VSP F400, F600, F800, setting Admin services EVS is not required.
- Create a storage pool

NAS Platform manages system drive volumes that are allocated to a file server and creates a storage pool by using required capacity from the system drives. The method depends on the version of NAS Platform.
  - NAS Platform v12.1 or later
    When DP pools are specified, Device Manager automatically creates storage pools that have suitable performance and expandability. Device Manager automatically allocates volumes and specifies the settings for writing to the system drives.
  - NAS Platform v11.3 or later
    Using Device Manager, create a storage pool by specifying the parity groups that have suitable performance and capacity for your storage pool purposes. Device Manager automatically allocates volumes and specifies the settings for writing to the system drives.
  - NAS Platform earlier than v11.3
    Create a volume on Device Manager and allocate the volumes to a file server. Then from Device Manager, open the SMU dialog box and create a storage pool. If you are using a version of NAS Platform earlier than v11.2, the system drives in the storage pool must belong to a system drive group.

- Create a file system and add file shares
  Create a file system on the storage pool, and add file shares. The method depends on the version of NAS Platform.
  - NAS Platform v10.2.3071 or later
    You can use Device Manager to create a file system and file shares at the same time.
  - NAS Platform earlier than v10.2.3071
    Launch SMU from Device Manager to create a file system and add file shares.

Tip: For SMU/NAS Manager v10.2.3071 or later, if configuration changes or version upgrades performed on the file server side are not applied to Device Manager, use SMU/NAS Manager to synchronize the information between the file server and Device Manager.

Related concepts

- About managing NAS Platform file systems and shares on page 246
- About registering and removing file servers on page 81
- About allocating volumes on page 201

Related tasks

- Creating a NAS Platform file system on page 264
- Adding NAS Platform file shares on page 270
- Creating a NAS Platform storage pool (v11.3 or later) on page 257
- Creating a NAS Platform storage pool (earlier than v11.3) on page 261

- Allocating volumes to selected file servers on page 206

Related references

- File share configuration examples on page 279
- Conditions for parity groups that are used in storage pools on page 253
- Conditions for DP pools that are used in storage pools on page 253

Conditions for parity groups that are used in storage pools

This topic describes the recommended conditions for parity groups from which you configure storage pools when creating a storage pool in NAS Platform.

To maximize the file server performance and efficiently use the storage capacity, use parity groups that meet the following conditions for creating or expanding storage pools:

- The parity groups are not used by any server or cluster other than the target server or cluster.
- A volume has not been created or only one volume has been created by using the entire capacity of the parity group.
- Parity groups are needed for each creation or expansion operation.
  - When NAS Platform version is earlier than 12.3:
    There are two or more parity groups.
  - When NAS Platform version is 12.3 or later:
    There are four or more parity groups.
- All parity groups have the same capacity.
- All parity groups have the same attribute for the following:
  - The system is an external storage system or internal storage system.
  - Drive type
  - Drive RPM
  - Chip Type
  - RAID level

If linking with NAS Platform v11.3 or later, when you select appropriate parity groups, Device Manager automatically creates or expands a storage pool with the optimum configuration based on best practices for NAS Platform.

Conditions for DP pools that are used in storage pools

This topic describes the recommended conditions for creating a storage pool in NAS Platform when using DP pools to create the system drives that comprise storage pools.

To maximize the file server performance and efficiently use the storage capacity, use DP pools that meet the following conditions for creating storage pools:

- The DP pool type is HDP.
- The DP pool is not being used in a server or cluster other than the target server or cluster.
- The DP volumes are not created from DP pools.
- The DP pool has the same attributes for the following:
  - The system is an external storage system or internal storage system.
  - Drive type
  - Drive RPM
  - Chip Type
  - RAID level
- For the Unified Storage 100 and Adaptable Modular Storage (AMS) 2000 family of storage systems, the stripe size of the DP pool must be 64 KB. The DP pool must be created by specifying the stripe size in Storage Navigator Modular 2.

If linking with NAS Platform v12.1 or later, when you select appropriate DP pools, Device Manager automatically creates or expands a storage pool with the optimum configuration based on best practices for NAS Platform.

Related references

- Best practices for configuring storage pools on page 254

Best practices for configuring storage pools

This topic describes the best practices for configuring storage pools to maximize the performance and reliability of NAS Platform and to efficiently use as much of the storage capacity as possible.

Storage pool configuration greatly affects performance, reliability, capacity efficiency, and expandability of file systems used. Follow the best practices suggested by NAS Platform for creating or expanding a storage pool. Following is an overview of the best practices.

DP pool and parity group

- Volume attributes
  Volumes that configure a storage pool must have the same attributes (drive type, drive speed in RPM, chip type, RAID level, and whether the drive is external or internal). If the attributes are different, overall performance is degraded by the volume with the lowest performance.

- Configuration of LUN paths
  If too many LUN paths exist between system drives and file server nodes, the system might stop. For NAS modules that are installed with VSP G400, G600, G800, and VSP F400, F600, F800, LUN paths are set automatically.

- Settings for writing to system drives
  You can achieve high I/O performance by specifying appropriate write settings for the performance of the device for each storage system type. In addition, if the stripe size of volumes can be defined, specify a size that is compatible with the characteristics of NAS Platform to improve data access efficiency. For NAS Platform v12.6 or later, NAS Platform sets the write settings automatically.

- Chunk size of storage pools
  For NAS Platform versions earlier than 11.2, you can specify the chunk size when creating a storage pool. There is a limit to the number of chunks that a file system can use. The chunk size affects the expandability of a file system. Small chunk sizes compromise expandability but have better performance. Large chunk sizes have better expandability but compromise performance.

DP pool only

- DP pool configuration
  For NAS Platform v12.1 or later, create the storage pool using an existing DP pool to support increased performance when expanding the storage pool. The storage pool and the DP pool must have a one-to-one relationship, meaning that a storage pool consists of a single DP pool. Also, use DP pools with uncreated DP volumes. DP volumes are used to automatically create the system drives for the storage pools and if DP volumes are already created, the capacity of the storage pool created automatically might decrease.
  To increase performance, configure:
  - The capacity and number of DP volumes for the storage pool as 8 TB x 8 DP volumes (total of 64 TB).
  - The number of volumes must be a minimum of four.
  - The number of drives that comprise a storage pool must be 24 or greater.

- DP pool stripe size
  For the Unified Storage 100 and Adaptable Modular Storage 2000 family of storage systems, the stripe size of the DP pool must be 64 KB.

Parity group only

- Number and capacity of volumes
  For creating or expanding a storage pool, separate out volumes from two or more parity groups to ensure performance.
  For using volumes for system drives, separate out fewer and larger volumes from each parity group. It is recommended to use the entire capacity of a parity group to create one volume. The reasons are:
  - Separating out many volumes from a parity group reduces performance.

- When volumes that constitute a storage pool are small, the lower the upper expansion limit becomes for a storage or file system.

In addition, the capacity of volumes that configure a storage pool needs to be equal. If you use volumes with different capacities, some areas might not be used and capacity efficiency deteriorates.

If linking with NAS Platform v11.3 or later, Device Manager automatically creates or expands a storage pool with the optimum configuration based on best practices. You just need to select the DP pools or parity groups that meet the conditions.

Related references

- Conditions for parity groups that are used in storage pools on page 253
- Conditions for DP pools that are used in storage pools on page 253

Notes on NAS Platform file system capacity

When creating a file system for NAS Platform family, the file system capacity entered by a user might differ from the capacity that is actually created.

Below are recommendations for specifying file system creation parameters, and reasons why specified capacity and actual capacity may differ.

For NAS Platform family, a file system is created in chunks. Therefore, the actual capacity of a file system is a multiple of the chunk size.

If the file server is NAS Platform v10.2.3071 or later, the value set for File System Capacity is rounded down to the nearest multiple of a chunk and displayed for Guideline Capacity as an approximation of the actual file system capacity. If you are creating a storage pool in Device Manager, the chunk size is set to 18 GB.

For example, if you specify 750 GB for File System Capacity, and then select a storage pool whose chunk size is 18 GB, 738 GB is displayed for Guideline Capacity.

When setting the file system capacity, consider the following recommendations:

- Create a file system with a capacity of 18 GB or more:
  If you set a value smaller than 18 GB for the file system capacity, the metadata area for managing the file system becomes large, and the actual file system capacity might become much larger than the specified value.
- If the chunk size is large for the file system capacity, the value of Guideline Capacity can easily differ from the actual file system capacity. On the other hand, if the chunk size is small, the maximum file system capacity also becomes small because the maximum number of chunks usable by the file system is limited.

If the free capacity of the selected storage system is smaller than the value specified for File System Capacity, a file system is created with the size of the free capacity of the storage pool. Therefore, the actual file system capacity differs from the value of Guideline Capacity.

Tip: If you are using a version of NAS Platform prior to v11.2, consider the following recommendations:

- Create a chunk size of 1 GB or more.
- The size of a created chunk may not match the specified size. If you create a file system by using a storage pool that has this chunk size, the capacity of the file system also may not match the value of **Guideline Capacity**.

For example, if you specify 10 GB for the chunk size when creating a storage pool, depending on the available capacity at the time of creation, a chunk might be created with 9.8 GB instead of 10 GB. In addition, when you expand the storage pool, depending on how much capacity you choose to expand, a chunk might then be created with a chunk size (for example 10.3 GB) that differs from the chunk size created originally (for example 9.8 GB). For this reason, the capacity of a file system you can create varies depending on the chunk that can be used when a file system is created.

Notes on setting LUN paths for NAS Platform

This topic describes recommended storage port guidelines when allocating volumes from storage systems to NAS Platform. For VSP G400, G600, G800 and VSP F400, F600, F800 with NAS modules, LUN paths are set automatically.

Recommendations for volumes and storage ports include:

- Set the number of volumes to less than or equal to the maximum values allowed for the number of storage ports that are being used.
- If the I/O load on the file server is high, increase the number of storage ports to be used according to the number of volumes.

For VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, VSP F1500, Virtual Storage Platform, Universal Storage Platform V/VM, or Unified Storage VM storage systems, the maximum number of volumes that can be allocated to a single file server (node) are:

- 64 x *number-of-storage-ports*

For mid-range storage systems:

- 16 x *number-of-storage-ports*

If the number of volumes exceeds the maximum value recommended, system performance might be affected.

Creating a NAS Platform storage pool (v11.3 or later)

You can create a storage pool on a cluster or a file server when the file server is NAS Platform v11.3 or later.

Note: When creating a storage pool, System LUs in NAS modules can not be specified. When expanding the storage pool, specify the number of system drives to be added as a number equal to or larger than the number of system drives specified when the storage pool was created.

Before you begin

Identify the target NAS Platform.

- For cluster configurations, identify the name of the target cluster.
- For non-cluster configurations, identify the name of the target file server.

Identify requirements for the storage pools to be created, and prepare necessary resources.

- For NAS Platform v12.3 or later, if you use only one parity group for creating a storage pool, create at least four volumes from that parity group in advance.
- For creating a storage pool, identify settings such as the storage pool name and capacity.
- To improve performance, reliability, and capacity efficiency of storage pools, set conditions for DP pools or parity groups that are used in storage pools.

Register the target NAS Platform to Device Manager.

- The file server must be registered in Device Manager.
- Device Manager must be linked to Admin services EVS, for file servers.
- The storage system with a NAS module must be registered in Device Manager.

Procedure

1. On the **Resources** tab, select **File Servers > All File Servers**. For NAS modules, on the **Resources** tab, select **Storage Systems > All Storage Systems**.
2. Select the target cluster or file server.
3. Click **Create Storage Pool** and specify the required items. For NAS Platform version 12.5 or later (excluding NAS modules), specify 7 and 68 for the **Host mode options** of the host group.
4. Click **Show Plan** and confirm that the information in the plan summary is correct. If changes are required, click **Back**.
5. (Optional) Update the task name and provide a description.
6. (Optional) Expand **Schedule** to specify the task schedule.

   You can schedule the task to run immediately or later. The default setting is **Now**. If the task is scheduled to run immediately, you can select **View task status** to monitor the task after it is submitted.
7. Click **Submit.**

   If the task is scheduled to run immediately, the process begins.

8. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.

Result

The storage pool is created on the NAS Platform file server.

Related concepts

- About managing NAS Platform file systems and shares on page 246

Related tasks

- Deleting NAS Platform storage pools on page 264

Related references

- Conditions for parity groups that are used in storage pools on page 253
- Conditions for DP pools that are used in storage pools on page 253

Expanding a NAS Platform storage pool (v11.3 or later)

If the file server is NAS Platform v11.3 or later and the capacity of a storage pool is insufficient, expand the already created storage pool.

Note: When expanding the storage pool, specify the number of system drives to be added as a number equal to or larger than the number of system drives specified when the storage pool was created.

Before you begin

Identify the target NAS Platform.

- For cluster configurations, identify the name of the target cluster.
- For non-cluster configurations, identify the name of the target file server.

Identify requirements for the storage pools to be expanded, and prepare necessary resources.

- For expanding a storage pool, identify settings such as the storage pool capacity.
- For DP volume configurations, the DP pools that comprise the storage pool must have enough capacity.
- To improve performance, reliability, and capacity efficiency of storage pools, set conditions for DP pools or parity groups that are used in storage pools.

Register the target NAS Platform to Device Manager.

- The file server must be registered in Device Manager.
- Device Manager must be linked to Admin services EVS.

Procedure

1. On the **Resources** tab, select **File Servers > All File Servers**. For NAS modules, on the **Resources** tab, select **Storage Systems > All Storage Systems**.
2. Select the target cluster or file server.
3. On the **Storage Pools** tab, select the row of the target storage pool.
4. Click **Expand Storage Pool** and specify the necessary items.
5. Click **Show Plan** and confirm that the information in the plan summary is correct. If changes are required, click **Back**.
6. (Optional) Update the task name and provide a description.
7. (Optional) Expand **Schedule** to specify the task schedule.

   You can schedule the task to run immediately or later. The default setting is **Now**. If the task is scheduled to run immediately, you can select **View task status** to monitor the task after it is submitted.
8. Click **Submit.**
   If the task is scheduled to run immediately, the process begins.
9. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.

Result

The storage pool capacity is expanded on the NAS Platform file server.

Related concepts

- About managing NAS Platform file systems and shares on page 246

Related references

- Conditions for parity groups that are used in storage pools on page 253
- Conditions for DP pools that are used in storage pools on page 253
- Notes on setting LUN paths for NAS Platform on page 257

Changing a NAS Platform module storage pool name

This topic describes changing the name of an existing storage pool when the file server is NAS Platform v11.3 or later.

Note: You can also change a storage pool name when the file server is a NAS Platform earlier than v11.3 by changing the NAS Platform storage pool settings.

Before you begin

- For cluster configurations, identify the name of the target cluster
- For non-cluster configurations, identify the name of the target file server
- For editing a storage pool, identify settings such as storage pool name
- The file server must be registered in Device Manager
- Device Manager must be linked to Admin services EVS

Procedure

1. On the **Resources** tab, select **File Servers > All File Servers**. For NAS modules, on the **Resources** tab, select **Storage Systems > All Storage Systems**.
2. Select the target cluster or file server.
3. On the **Storage Pools** tab, select the row of the target storage pool.
4. Click **Edit Storage Pool** and specify the necessary items.
5. (Optional) Update the task name and provide a description.
6. Click **Submit**.
7. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.

Result

The storage pool name is changed.

Related tasks

- Changing NAS Platform storage pool settings on page 263

Creating a NAS Platform storage pool (earlier than v11.3)

When the file server is NAS Platform versions earlier than 11.3, launch System Management Unit (SMU) from Device Manager to create a storage pool on the file server.

Before you begin

- The file server must be registered in Device Manager
- Allocate volumes to the cluster or file server
- For cluster configurations, identify the name of the target cluster
- For non-cluster configurations, identify the name of the target file server
- For creating a storage pool, identify settings such as the storage pool name and capacity

Procedure

1. On the **Resources** tab, select **File Servers > All File Servers**.
2. Expand the tree and select the target cluster or file server.
3. From the **Actions** menu, select **Manage System Drives**.
4. In the **SMU** dialog box, change the settings so that the required system drive can be accessed.

   On the **System Drives** tab of Device Manager, **Allowed** is displayed in the **Access** column of the target system drive. If you are using a version of NAS Platform earlier than v11.2, continue to step 5. If you are using v11.2 or a later version, skip to step 7.
5. From the **Actions** menu, select **Create System Drive Group**. If a system drive that belongs to the same parity group or DP pool as for the

target system drive is included in the already existing system drive group, select **Manage System Drive Groups** and add the target system drive into that system drive group.

6. In the **SMU** dialog box, perform the operations.
7. On the **Storage Pools** tab, click **Create Storage Pool**.
8. In the **SMU** dialog box, specify the required items.

   For explanations of items and procedures in the **SMU** dialog box, see the SMU online help.
9. Check the created storage pool in Device Manager, on the **Storage Pools** tab of the target cluster or file server.

Result

The storage pool is created on the NAS Platform file server.

Expanding a NAS Platform storage pool (earlier than v11.3)

When the file server is NAS Platform versions earlier than 11.3, if the capacity of a storage pool server is insufficient, launch System Management Unit (SMU) from Device Manager to expand the already created storage pool on the file server.

Before you begin

- The file server must be registered in Device Manager
- Allocate volumes to the cluster or file server
- For cluster configurations, identify the name of the target cluster
- For non-cluster configurations, identify the name of the target file server
- For expanding a storage pool, identify settings such as the storage pool name and capacity

Procedure

1. On the **Resources** tab, select **File Servers > All File Servers**.
2. Expand the tree and select the target cluster or file server.
3. From the **Actions** menu, select **Manage System Drives**.
4. In the **SMU** dialog box, change the settings so that the required system drive can be accessed.

   On the **System Drives** tab of Device Manager, **Allowed** is displayed in the **Access** column of the target system drive. If you are using a version of NAS Platform earlier than v11.2, continue to step 5. If you are using v11.2 or a later version, skip to step 7.
5. From the **Actions** menu, select **Create System Drive Group**. If a system drive that belongs to the same parity group or DP pool as for the target system drive is included in the already existing system drive group, select **Manage System Drive Groups** and add the target system drive into that system drive group.
6. In the **SMU** dialog box, specify the required items.

7. On the **Storage Pools** tab, click **Manage Storage Pool**.
8. In the **SMU** dialog box, specify the required items.

   For explanations of items and procedures in the **SMU** dialog box, see the SMU online help.
9. Check the expanded storage pool in Device Manager, in the **Storage Pools** tab of the target cluster or file server.

Result

The storage pool is expanded on the NAS Platform file server.

Changing NAS Platform storage pool settings

When the file server is NAS Platform, launch System Management Unit (SMU)/NAS Manager from Device Manager to change settings of a storage pool, such as accessibility to the storage pool, and availability of auto-expansion for the file system. For NAS Platform versions earlier than 11.3, use this procedure to change the storage pool name.

Before you begin

- The file server must be registered in Device Manager
- For cluster configurations, identify the name of the target cluster
- For non-cluster configurations, identify the name of the target file server
- For changing a storage pool settings, identify the storage pool name and settings to be changed

Procedure

1. On the **Resources** tab, select **File Servers > All File Servers**. For NAS modules, on the **Resources** tab, select **Storage Systems > All Storage Systems**.
2. Expand the tree and select the target cluster or file server.
3. On the **Storage Pools** tab, select the target storage pool.
4. Click **Manage Storage Pool.**
5. In the **SMU/NAS Manager** dialog box, specify the required items.
6. On the **Storage Pools** tab, click **Manage Storage Pool**.
7. In the **SMU/NAS Manager** dialog box, specify the required items.

   For explanations of items and procedures in the **SMU/NAS Manager** dialog box, see the SMU/NAS Manager online help.
8. Check the changed storage pool settings in Device Manager, on the **Storage Pools** tab of the target cluster or file server.

Result

The storage pool settings are changed on the NAS Platform file server.

Deleting NAS Platform storage pools

You can delete storage pools on a cluster or a file server when the storage pools are no longer needed.

If both a storage pool and a LUN path are deleted, and:

- The storage pool consists of DP pools, the DP volumes are also deleted.
- The storage pool consists of basic volumes, the basic volumes are also deleted.

Before you begin

Identify the name of the target storage pool.

Procedure

1. On the **Resources** tab, select **File Servers > All File Servers**. For NAS modules, on the **Resources** tab, select **Storage Systems > All Storage Systems**.
2. Select the target cluster or file server.
3. On the **Storage Pools** tab, select the storage pool to delete. Volumes related to the target storage pools are also deleted.
   - NAS Platform v11.3 or later:
     Click **Delete Storage Pools** and skip to step 4.
   - NAS Platform versions earlier than 11.3:
     Click **Manage Storage Pool**. In the **SMU** dialog box, delete the storage pool, and then skip to step 6.
4. Click **Submit.**
   If the task is scheduled to run immediately, the process begins.
5. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.
6. On the **Storage Pools** tab of the target cluster or file server, verify the deleted storage pool.

Result

The storage pool is deleted on the NAS Platform file server.

Creating a NAS Platform file system

Depending on the version of NAS Platform you are using, create a file system and file shares on the file server by using either HCS or System Management Unit (SMU)/NAS Manager.

Before you begin

- The file server must be registered in Device Manager
- Create a storage pool
- For cluster configurations, identify the target cluster name

- For non-cluster configurations, identify the target file server name
- Identify the EVS to which the file system should be assigned.
- Identify the storage pool name
- Identify file system name, allocation type, and capacity
- Identify block size
- In addition, verify the following when using NAS Platform version 10.2.3071 or later:
  - Device Manager must be linked to Admin services EVS
  - Identify file share protocols, names, and paths
  - When intending to set access permissions for CIFS shares, add a CIFS server from the CIFS setup window in System Management Unit, and set up information on an authentication server, such as Active Directory.

Procedure

1. On the **Resources** tab, select **File Servers > All File Servers**. For NAS modules, on the **Resources** tab, select **Storage Systems > All Storage Systems**.
2. From the Servers/Clusters list, select the row of the target cluster or file server, and click **Create File System**.
3. Create a storage pool or file share depending on the version of NAS Platform:
   - NAS Platform v10.2.3071 or later:
     Specify the required items and skip to step 4.
   - NAS Platform versions earlier than 10.2.3071:
     In the **SMU** dialog box, select whether to create a new storage pool or use an already created storage pool, specify the required items, and then skip to step 7.
4. Click **Show Plan** and confirm that the information in the plan summary is correct. If changes are required, click **Back**.
5. Click **Submit.**
   If the task is scheduled to run immediately, the process begins.
6. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.
7. To verify that the file system was created, select the target cluster or file server, and then click the **File Systems** tab. To verify that the file share was created, select the target file system, and then click the **Shares/ Exports** tab.

Result

The file system or file shares are created on the file server.

Related concepts

- About managing NAS Platform file systems and shares on page 246

Related tasks

- Changing NAS Platform file share settings on page 271
- Deleting a NAS Platform file system on page 269

Related references

- Workflow for providing NAS Platform file shares on page 250
- Notes on NAS Platform file system capacity on page 256

Expanding a NAS Platform file system

Depending on the version of NAS Platform you are using, increase the maximum allowed file system size by using either HCS or System Management Unit (SMU)/NAS Manager.

Before you begin

- The file server must be registered in Device Manager
- For cluster configurations, identify the target cluster name
- For non-cluster configurations, identify the target file server name
- Identify the name of the target file system
- Identify the maximum size of the capacity to be allocated
- In addition, verify the following when using NAS Platform v11.1 or later:
  - Device Manager must be linked to Admin services EVS

Procedure

1. On the **Resources** tab, select **File Servers > All File Servers**. For NAS modules, on the **Resources** tab, select **Storage Systems > All Storage Systems**.
2. From the tree, select the target cluster or file server, and on the **File Systems** tab, select the file system to expand.
3. Expand the file system depending on the version of NAS Platform:
   - NAS Platform v12.0 or later:
     Click **Expand File System**, enter the capacity, and then skip to step 4.
   - NAS Platform v11.1 or later and earlier than v12.0:
     For On Demand Allocation, click **Expand File System**, specify the allocation limit size, and then skip to step 4. For Full Allocation, click **Manage File System**, and in the **SMU** dialog box, expand the file system. Then skip to step 7.
   - NAS Platform versions earlier than 11.1:
     Click **Manage File System**. In the **SMU** dialog box, expand the file system, and then skip to step 7.
4. Click **Show Plan** and confirm that the information in the plan summary is correct. If changes are required, click **Back**.
5. Click **Submit.**
   If the task is scheduled to run immediately, the process begins.

6. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.
7. On the **File Systems** tab of the target cluster or file server, verify the expanded file system.

Result

The selected file system is expanded.

Related concepts

- About managing NAS Platform file systems and shares on page 246

Changing NAS Platform file system settings

When the file server is NAS Platform v10.1.3070 or later, launch System Management Unit (SMU)/NAS Manager from Device Manager to change a file system's name, or change the threshold value for the usage rate.

Before you begin

- The file server must be registered in Device Manager
- Identify the name of the target cluster or target file server.

Procedure

1. On the **Resources** tab, select **File Servers > All File Servers**. For NAS modules, on the **Resources** tab, select **Storage Systems > All Storage Systems**.
2. In the tree, select the target cluster or file server, and on the **File Systems** tab, select the file system.
3. Click **Manage File System**.
4. In the **SMU/NAS Manager** dialog box, specify the required items.

   For explanations of items and procedures in the **SMU/NAS Manager** dialog box, see the SMU/NAS Manager online help.
5. Select the target cluster or file server in Device Manager, and on the **File Systems** tab, confirm that the settings are changed.

Result

The file system settings are changed.

Related concepts

- About managing NAS Platform file systems and shares on page 246

Related tasks

- Creating a NAS Platform file system on page 264

Related references

- Workflow for providing NAS Platform file shares on page 250
- Notes on NAS Platform file system capacity on page 256

Mounting NAS Platform file systems

Depending on the version of NAS Platform you are using, mount file systems by using either HCS or System Management Unit (SMU)/NAS Manager.

Before you begin

- The file server must be registered in Device Manager
- For cluster configurations, identify the target cluster name
- For non-cluster configurations, identify the target file server name
- Identify the name of the target file system
- In addition, verify the following when using NAS Platform v11.1 or later:
  - Device Manager must be linked to Admin services EVS

Procedure

1. On the **Resources** tab, select **File Servers > All File Servers**. For NAS modules, on the **Resources** tab, select **Storage Systems > All Storage Systems**.
2. From the tree, select the target cluster or file server, and on the **File Systems** tab, select the file systems to mount.
3. Mount the file systems depending on the version of NAS Platform:
   - NAS Platform v11.1 or later:
     Click **Mount File Systems** and skip to step 4.
   - NAS Platform versions earlier than 11.1:
     Click **Manage File System**. In the **SMU** dialog box, mount the file systems, and then skip to step 6.
4. Click **Submit.**
   If the task is scheduled to run immediately, the process begins.
5. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.
6. On the **File Systems** tab of the target cluster or file server, verify the mounted file systems.

Result

The selected file systems are mounted.

Unmounting NAS Platform file systems

Depending on the version of NAS Platform you are using, unmount file systems by using either HCS or System Management Unit (SMU)/NAS Manager.

Before you begin

- The file server must be registered in Device Manager
- For cluster configurations, identify the target cluster name

- For non-cluster configurations, identify the target file server name
- Identify the name of the target file system
- In addition, verify the following when using NAS Platform v11.1 or later:
  - Device Manager must be linked to Admin services EVS

Procedure

1. On the **Resources** tab, select **File Servers > All File Servers**. For NAS modules, on the **Resources** tab, select **Storage Systems > All Storage Systems**.
2. From the tree, select the target cluster or file server, and on the **File Systems** tab, select the file systems to unmount.
3. Unmount the file systems depending on the version of NAS Platform:
   - NAS Platform v11.1 or later:
     Click **Unmount File Systems** and skip to step 4.
   - NAS Platform versions earlier than 11.1:
     Click **Manage File System**, in the **SMU** dialog box, unmount the file systems, and then skip to step 6.
4. Click **Submit.**
   If the task is scheduled to run immediately, the process begins.
5. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.
6. On the **File Systems** tab of the target cluster or file server, verify the unmounted file systems.

Result

The selected file systems are unmounted.

Deleting a NAS Platform file system

If a NAS Platform file system is no longer needed, for example if the file system is being decommissioned and the needed data is saved, you can delete the file system.

Before you begin

Identify the name of the target file system.

Procedure

1. On the **Resources** tab, select **File Servers > All File Servers**. For NAS modules, on the **Resources** tab, select **Storage Systems > All Storage Systems**.
2. From the list of file systems, select the file system that you want to delete.
   - NAS Platform v11.1 or later:
     Click **Delete File Systems** and skip to step 3.

- NAS Platform versions earlier than 11.1:
  Click **Manage File System**. In the **SMU** dialog box, delete the file system, and then skip to step 5.

3. Click **Submit.**
   If the task is scheduled to run immediately, the process begins.
4. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.
5. From the **All File Servers** list of file systems, verify the deleted file system. For NAS modules, from the **All Storage Systems** list, verify the deleted file system.

Result

The file system you deleted no longer appears in the list of file servers/NAS modules.

Adding NAS Platform file shares

Depending on the version of NAS Platform you are using, you can add file shares to existing file systems on the file server by using either HCS or System Management Unit (SMU)/NAS Manager.

Before you begin

- The file server must be registered in Device Manager
- For cluster configurations, identify the target cluster name
- For non-cluster configurations, identify the target file server name
- Identify the EVS to which the target file system belongs
- Identify the name of the target file system
- Identify file share protocols, names, paths, and access permissions
- In addition, verify the following when using NAS Platform v11.1 or later:
  - Device Manager must be linked to Admin services EVS
  - When intending to set access permissions for CIFS shares, add a CIFS server from the CIFS setup window in System Management Unit, and set up information on an authentication server (such as Active Directory).

Procedure

1. On the **Resources** tab, select **File Servers > All File Servers**. For NAS modules, on the **Resources** tab, select **Storage Systems > All Storage Systems**.
2. From the tree, select the target cluster or file server, and on the **File Systems** tab, select the file system where file shares will be added.
3. Add the file shares depending on the version of NAS Platform Platform:
   - NAS Platform v11.1 or later:
     Click **Add Share/Export**, specify the required items, and skip to step 4.

- NAS Platform versions earlier than 11.1:
  Click **Add CIFS Share** or **Add NFS Export**, depending on the protocol of the file share you are adding. In the **SMU** dialog box, add the file share, and skip to step 7.

4. Click **Show Plan** and confirm that the information in the plan summary is correct. If changes are required, click **Back**.
5. Click **Submit.**
   If the task is scheduled to run immediately, the process begins.
6. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.
7. On the **Shares/Exports** tab of the EVS where the target file system belongs, verify the added file shares.

Result

The file shares are added to existing the file systems on the file server.

Related tasks

- Changing NAS Platform file share settings on page 271
- Deleting NAS Platform file shares on page 272

Related references

- Workflow for providing NAS Platform file shares on page 250
- File share configuration examples on page 279

Changing NAS Platform file share settings

Depending on the version of NAS Platform you are using, edit settings for existing file shares using either HCS or System Management Unit (SMU)/NAS Manager.

Note: If the NAS Platform version is 12.5.4025.00 or later and earlier than version 12.7.4221.05, then Edit CIFS Share cannot be used to set file shares as root directories. To set a file share as a root directory, use SMU/NAS Manager or the NAS Platform CLI.

Before you begin

- The file server must be registered in Device Manager
- Identify the file share name
- For cluster configurations, identify the target cluster name of the file share
- For non-cluster configurations, identify the target file server name of the file share
- Identify the EVS to which the file share belongs
- In addition, verify the following when using NAS Platform v11.1 or later:
  - Device Manager must be linked to Admin services EVS
  - When intending to set access permissions for CIFS shares, add a CIFS server from the CIFS setup window in System Management Unit, and

set up information on an authentication server (such as Active Directory).

Procedure

1. On the **Resources** tab, select **File Servers > All File Servers**. For NAS modules, on the **Resources** tab, select **Storage Systems > All Storage Systems**.
2. From the tree, select the target cluster or file server.
3. From the tree, select the target EVS, and on the **Shares/Exports** tab, select the file share you want to modify.
4. Change the file share settings depending on the version of NAS Platform:
   - NAS Platform v11.1 or later:
     Click **Edit Share/Export**, specify the required items, and skip to step 5.
   - NAS Platform versions earlier than 11.1:
     From the list of file shares, select the file share you want to modify, and click **Manage Share/Export**. In the **SMU** dialog box, change the file share, and skip to step 8.
5. Click **Show Plan** and confirm that the information in the plan summary is correct. If changes are required, click **Back**.
6. Click **Submit.**
   If the task is scheduled to run immediately, the process begins.
7. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.
8. On the **Shares/Exports** tab of the target EVS, verify the modified file share settings. In SMU/NAS Manager, click **Manage Share/Export** to verify the modified file share settings.

Result

The file share settings are changed.

Related tasks

- Adding NAS Platform file shares on page 270

Deleting NAS Platform file shares

If NAS Platform file shares are no longer needed, you can delete the file shares.

Before you begin

Identify the name of the target file share.

Procedure

1. On the **Resources** tab, select **File Servers > All File Servers**. For NAS modules, on the **Resources** tab, select **Storage Systems > All Storage Systems**.
2. From the list of file systems, select the file system that has the file shares you want to delete.
3. On the **Shares/Exports** tab, select the target file share.
   - NAS Platform v11.1 or later:
     Click **Delete Shares/Exports**, and then skip to step 4.
   - NAS Platform versions earlier than 11.1:
     Click **Manage Share/Export**. In the **SMU** dialog box, delete the file share, and then skip to step 6.
4. Click **Submit.**
   If the task is scheduled to run immediately, the process begins.
5. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.
6. On the **Shares/Exports** tab, verify the deleted file shares.

Result

The file shares you deleted no longer appear in the list of file shares.

Managing NAS Platform F and Data Ingestor file systems and shares

This module describes how to create file systems and shares for NAS Platform and Data Ingestor file servers.

About managing NAS Platform F and Data Ingestor file systems and shares

When introducing a new NAS Platform F or Data Ingestor file server, register the file server in Device Manager by using File Services Manager management software, then allocate storage system volumes to the file server to create file systems and shares.

Data Ingestor tasks can be performed using File Services Manager. Common tasks can also be performed from the HCS GUI, for ease of use.

Depending on the File Services Manager software version, and the Hitachi Command Suite version in use, tasks will either describe using File Services Manager, or describe using the Hitachi Command Suite GUI.

If File Services Manager v3.1 or later is installed on the same server as Device Manager you can:

- Create a file system

You can create a file system on a file server by using storage system volumes. You can also allocate volumes to the file server and create file shares at the same time you create a file system.

- Expand a file system
  You can use Top 10 File Systems from the dashboard to check for file systems with a high capacity usage. If a file system does not have sufficient free capacity, you can allocate the necessary amount of volumes to expand the file system. When settings are specified so that SNMP traps are received from the file server, you can check, by using alerts, whether the usage rate of a file system exceeds a threshold. For details about settings for receiving SNMP traps, see the *Hitachi Command Suite Administrator Guide*.
- Delete a file system
  You can delete file systems that are no longer required. Volumes used by those file systems can be reused for other file systems. If you do not reuse volumes for new file systems, consider releasing the volume allocation from the file server, after deleting the file systems, so the volume can be used elsewhere.
- Add file shares
  You can add file shares in response to a request from a file server user. Check the purpose of the file shares or the free capacity of the file systems, and then select a file system to which file shares are to be added.
- Edit or release file shares
  You can edit file share attributes and release unnecessary file shares. When you want to release shares that are used by both CIFS and NFS protocols, releasing one of the protocols allows both protocols to be released.

Tip: To check or change the capacity of file shares, start File Services Manager and then perform the operation. You can start File Services Manager from the Device Manager file server view.

Related concepts

- Allocating storage on page 141
- About registering and removing file servers on page 81

Related references

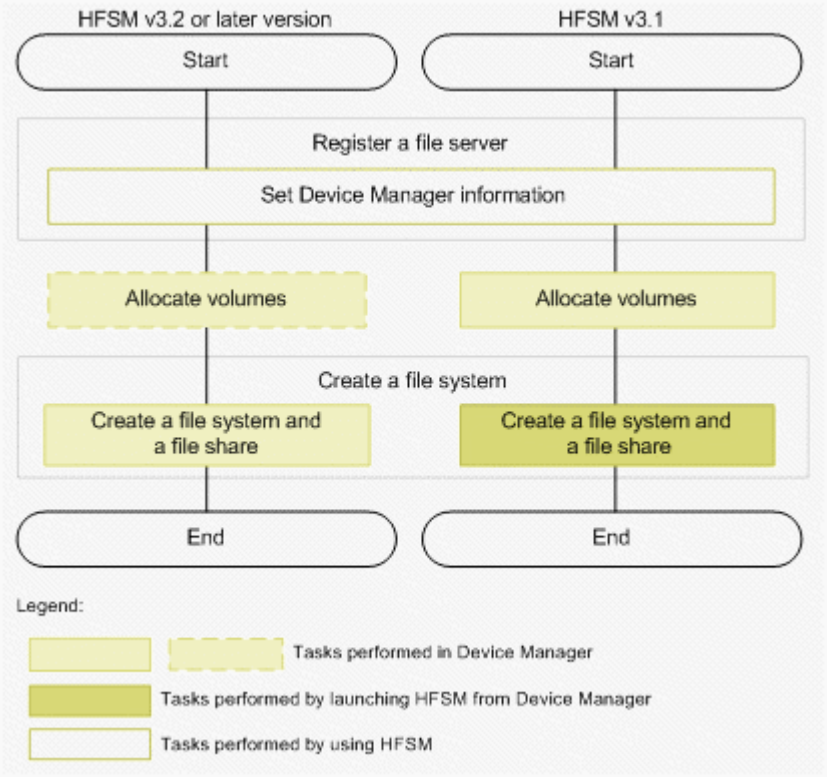
- Workflow for providing HNAS F and HDI file shares (HFSM) on page 275
- File share configuration examples on page 279

Workflow for providing HNAS F and HDI file shares (HFSM)

If File Services Manager v3.1 or later and Device Manager are installed on the same server, you can create file systems and file shares from the Device Manager window for File Services Manager F or Data Ingestor file servers.

Tip: If Device Manager and File Services Manager are installed on different servers, or if the version of a linked File Services Manager is earlier than v3.1, use Device Manager to allocate volumes to the file server before using File Services Manager.

The workflow for registering a file server to provide file systems and shares is indicated in the following figure:



- Register a file server
  By setting Device Manager information in File Services Manager, file server information managed in File Services Manager is sent to Device Manager to register file servers to be managed by Device Manager.
- Allocate volumes use Device Manager to allocate volumes to a file server. If File Services Manager v3.2 or a later version and Device Manager are installed on the same server, volumes that match the requirements can be automatically allocated when the file system is created.
- Create a file system
  Create a file system and a file share. The creation procedure varies depending on the version of File Services Manager, which is installed on the same server as Device Manager.

- File Services Manager v3.2 or later
  You can use Device Manager to create a file system. You can also allocate volumes to the file server and create file shares in a batch operation at the same time you create a file system. When the capacity of the file system to be created is specified, a volume allocation plan is automatically provided. A volume that is already allocated to the file server can also be used. If allocation is performed by specifying one file server (node) in a cluster, volumes are also automatically allocated to the other file server (node).
- File Services Manager v3.1
  You can create a file system and a file share by starting the File Services Manager dialog box from the Device Manager window.

Related concepts

- About managing NAS Platform F and Data Ingestor file systems and shares on page 273

Related tasks

- Creating HNAS F or HDI file systems (HFSM v3.2 or later) on page 276
- Creating HDI file systems (HFSM v3.1) on page 277
- Changing HNAS F or HDI file system settings (HFSM v3.1 or later) on page 278
- Registering file servers on page 84
- Allocating volumes to selected file servers on page 206

Related references

- File share configuration examples on page 279

Creating HNAS F or HDI file systems (HFSM v3.2 or later)

If File Services Manager v3.2 or later is installed on the same server as Device Manager, you can allocate volumes to NAS Platform F or Data Ingestor file servers, create a file system, and create file shares in a batch operation. You can also use volumes already allocated to the file server to create a file system.

Note: To use a volume that has already been allocated, alternate paths belonging to the same host group are set between the file server and volume.

Before you begin

- Identify the name of the target cluster or target file server.
- Identify the name and capacity of the file systems to be created.
- Identify the file share protocol if creating shares.

Procedure

1. On the **Resources** tab, from the tree, select **File Servers**, and then select **All File Servers**.
2. From the Servers/Clusters list, select the row of the target cluster, and then click **Create File System**. If you will use volumes that have already been allocated to a file server, select the cluster of the target file server (node) in the tree view. Select the target volumes on the **Volumes** tab, click **Create File System**, and specify the required items.
3. Click **Show Plan** and confirm that the information in the plan summary is correct. If changes are required, click **Back**.
4. Click **Submit.**
   If the task is scheduled to run immediately, the process begins.
5. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.
6. Select the target cluster, and on the **File Systems** tab, verify the created file system. Select the target file system, and then on the **Shares** tab, verify the created file share.

Result

The file system and file shares are created on the file server.

For details about the requirements for the environment settings of the file server, see the *Hitachi Command Suite Administrator Guide*.

Related tasks

- Changing HNAS F or HDI file system settings (HFSM v3.1 or later) on page 278
- Deleting HNAS F or HDI file systems on page 279

Related references

- Workflow for providing HNAS F and HDI file shares (HFSM) on page 275
- File share configuration examples on page 279
- Notes on performing quick formats on page 151

Creating HDI file systems (HFSM v3.1)

When File Services Manager v3.1 is installed on the same server as Device Manager, you can open an File Services Manager dialog box from Device Manager, and perform operations such as creating file systems on an Data Ingestor file server. You can also create file shares at the same time you create a file system.

Before you begin

- Identify the name of the target cluster or target file server.
- Allocate volumes to the file server.
- Identify the name and capacity of the file systems to be created.

- Identify the name, protocol, and owners of file shares.

Procedure

1. On the **Resources** tab, from the tree, select **File Servers**, then select **All File Servers**.
2. From the Servers/Clusters list of file servers, select the row of the cluster of the target file server (node).
3. Click **Create File System** or **Create and Share File System**.
4. From the **HFSM** dialog box, specify the required items.

   For explanations of items and procedures in the **HFSM** dialog box, see the HFSM online help.
5. Select the target cluster in Device Manager, and on the **File Systems** tab, verify the created file system. Select the target file system, and on the **Shares** tab, verify the created file share.

Result

The Data Ingestor file system is created on the file server.

For details about the requirements for the environment settings of the file server, see the *Hitachi Command Suite Administrator Guide*.

Related tasks

- Changing HNAS F or HDI file system settings (HFSM v3.1 or later) on page 278
- Deleting HNAS F or HDI file systems on page 279

Related references

- Workflow for providing HNAS F and HDI file shares (HFSM) on page 275
- File share configuration examples on page 279

Changing HNAS F or HDI file system settings (HFSM v3.1 or later)

When File Services Manager v3.1 or later is installed on the same server as Device Manager, you can open an File Services Manager dialog box from Device Manager, and perform operations such as expanding a file system and adding file shares created on an NAS Platform F or Data Ingestor file server.

Procedure

1. On the **Resources** tab, from the tree, select **File Servers**.
2. In the tree, select the cluster of the target file server (node), and then select the row of the target file systems in the **File Systems** tab.
3. Depending on the task that you are doing, click **Expand File System** or **Add Share**.
4. In the **HFSM** dialog box, specify the required items.

   For explanations of items and procedures in the **HFSM** dialog box, see the HFSM online help.

5. Select the target cluster in Device Manager, and on the **File Systems** tab, verify the expanded file system. Select the target file system, and on the **Shares** tab, verify the added file share.

Result

The file system settings are changed.

Related tasks

- Creating HDI file systems (HFSM v3.1) on page 277
- Creating HNAS F or HDI file systems (HFSM v3.2 or later) on page 276

Related references

- Workflow for providing HNAS F and HDI file shares (HFSM) on page 275

Deleting HNAS F or HDI file systems

You can delete HNAS F or HDI file systems when they are no longer needed. Save required data prior to deleting the file system.

Before you begin

Identify the name of the target file system.

Procedure

1. On the **Resources** tab, from the tree, select **File Servers**, and then select **All File Servers**.
2. From the list of file systems, select the file system that you want to delete.

   You can delete the file systems of File Services Manager v3.2 or later and File Services Manager v3.1.
3. Click **Delete File Systems**.
4. Click **Show Plan** and confirm that the information in the plan summary is correct. If changes are required, click **Back**.
5. Click **Submit.**
   If the task is scheduled to run immediately, the process begins.
6. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.
7. From the **All File Servers** list of file systems, verify the deleted file system.

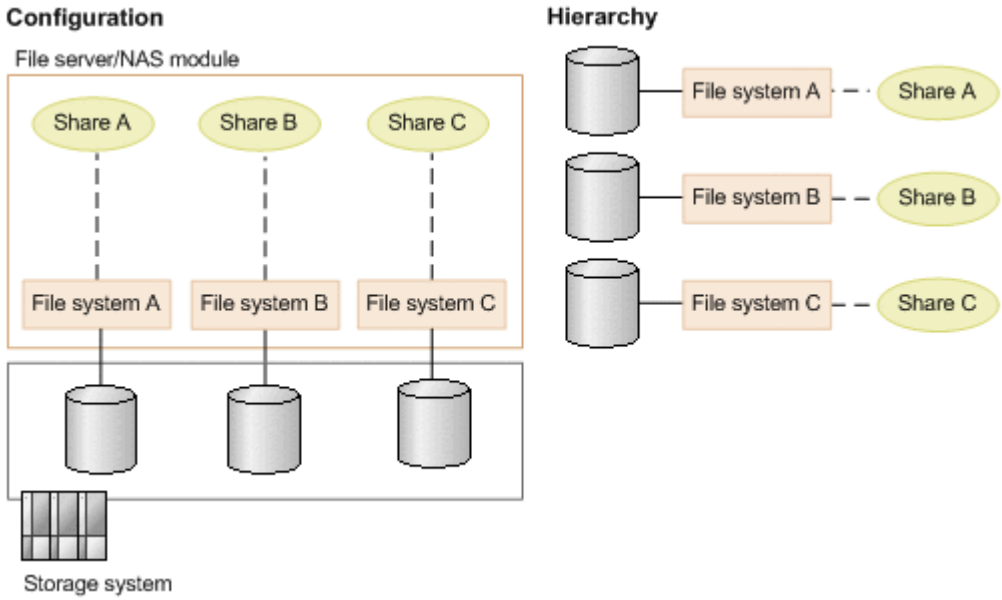
Result

The file system you deleted no longer appears in the list of file servers.

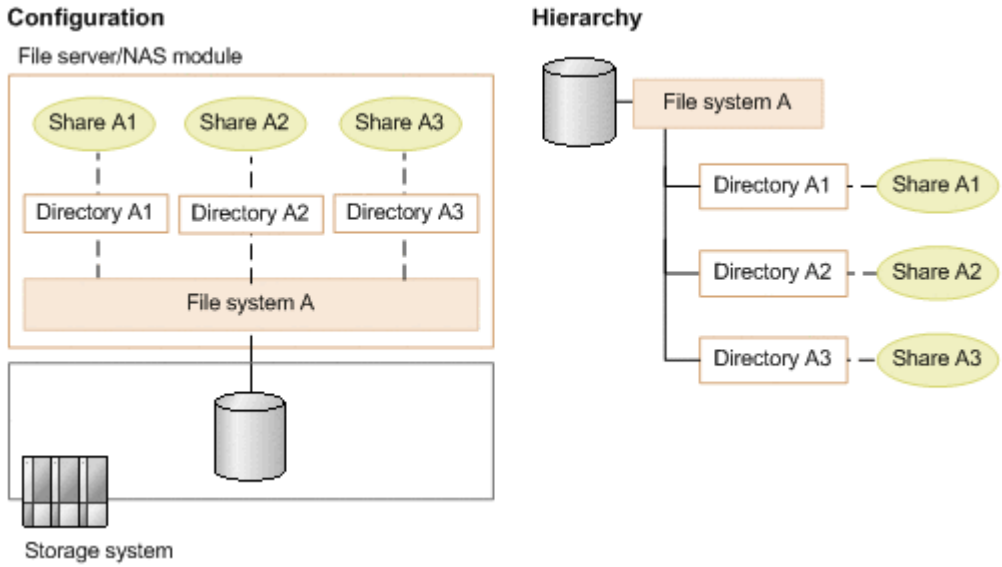
File share configuration examples

File systems can be shared in several ways, for example, sharing file system root directories, or sharing specific file system sub-directories.

In the illustration below, multiple root file systems have been shared (specified as an export point).



In the illustration below, multiple file system sub-directories have been created using a single file system, and shared (specified as an export point).



Related concepts

- About managing NAS Platform file systems and shares on page 246
- About managing NAS Platform F and Data Ingestor file systems and shares on page 273

Related tasks

- Adding NAS Platform file shares on page 270
- Creating HNAS F or HDI file systems (HFSM v3.2 or later) on page 276

- Creating HDI file systems (HFSM v3.1) on page 277

Related references

- Workflow for providing NAS Platform file shares on page 250
- Workflow for providing HNAS F and HDI file shares (HFSM) on page 275

Managing copy pairs

You can use the Replication Manager to manage copy pairs.

About pair management

Replication Manager is used to manage copy pairs, including remote replication. HCS provides multiple pair management solutions based on Replication Manager licensing:

- Full Replication Manager functionality is available from the Tools menu (HRpM license required). You can use wizards to visualize complex replication configurations while defining, editing, viewing, and troubleshooting copy pairs. For information about setting up and using Replication Manager, see the *Hitachi Command Suite Replication Manager User Guide*.
- A basic version of Replication Manager (included with the Device Manager license) that provides for common pair management tasks (for example, create a pair) using the Actions menu, Manage Replication. See HCS copy pair management operations (no Replication Manager license) on page 281 for a summary of the basic capabilities.
- The Replication tab supports a convenient set of remote copy configurations, allowing the storage administrator to perform operations that would otherwise require accessing the Replication Manager GUI. (HRpM license required.) The Replication tab also allows you to configure and monitor alerts, check the status of a replication configuration, and resolve errors. See HCS Replication tab operations (with Replication Manager license) on page 288 for a summary of capabilities.

HCS copy pair management operations (no Replication Manager license)

You can perform the following pair management functions in Hitachi Command Suite without a Replication Manager license:

- Define, create, and delete open volume copy pairs
- Modify copy-pair tasks
  - create: Creates a copy pair
  - split: Separates a copy pair
  - resync: Synchronizes a secondary volume with the primary volume
  - restore: Restores the copy pair from a secondary volume to the primary volume

    - delete: Releases a copy pair and changes the status to simplex
- View information from the Hosts, Storage Systems, and Copy-Pair Configuration Definitions trees
- View tasks
- View workflows
- Add, edit, or delete command devices
- View event logs
- Switch to maintenance mode

To access other pair management functionality (whether through the Replication Manager GUI or the Replication tab), you must register a Replication Manager license and log in to HCS with Replication Manager permissions. In addition, requisite copy types (such as ShadowImage and TrueCopy), CCI, and Device Manager agent must be installed and configured. For requirements and procedures, refer to the release notes and the *Hitachi Command Suite Administrator Guide*.

Defining copy pairs

Before you define copy pairs, determine the volumes that will be used to define copy pairs and verify that the volumes are allocated to a host.

Procedure

1. From the **Actions** menu, select **Manage Replication** to start Replication Manager.
2. Display a list of volumes from the **Hosts** tree or the **Storage Systems** tree.
3. On the **Unpaired** tab, select the volumes from which copy pairs will be defined and click **Pair Management**.

   The Pair Configuration wizard opens.
4. Follow the steps in the Pair Configuration wizard to define copy pairs.

   For details about how to specify copy-pair conditions and how to use the Pair Configuration wizard, see the Replication Manager online help.
5. From the **Explorer** menu, select **Tasks**, and then select **Tasks** again to confirm that the task was properly executed.
6. Confirm the copy-pair definitions in the **Pair Configurations** or **Storage Systems** view in Replication Manager.

   From the list of volumes associated with **Hosts** tree or **Storage Systems**, view the current copy pair status.



Tip: To display the most recent SplitTime in Device Manager after performing operations on a Copy-on-Write Snapshot or Thin Image copy pair, you need to refresh the storage system information.

Changing the status of a copy pair

You can change the status of a copy pair.

Procedure

1. From the **Actions** menu, select **Manage Replication** to start Replication Manager.
2. From the **Hosts** tree, select a host and click **Refresh Hosts**.
   All copy pairs associated with that host are refreshed.
3. Display the list of volumes from the **Hosts** tree.
4. Click the icon to change the pair status.
5. Change the copy pair status by using the **Change Pair Status** wizard.
   For details on the conditions under which certain copy pair statuses can be changed and on how to change the status of a copy pair by using the Change Pair Status wizard, see the Replication Manager online Help.
6. From the **Explorer** menu, select **Tasks**, then select **Tasks** again to confirm that the task is completed.
7. Confirm the copy pair status was changed by viewing the list of volumes in the **Hosts** view or the **Storage** view in Replication Manager.
   From the list of volumes in the **Hosts** tree or the **Storage Systems** tree, view the current copy pair status.

Tip: To display the most recent SplitTime in Device Manager after performing operations on a Copy-on-Write Snapshot or Thin Image copy pair, you need to refresh the storage system information.

Managing command devices

Command devices are necessary for CCI to communicate with storage systems.

You can use the Replication Manager to manage command devices. For more information on pair management using Replication Manager, see the *Hitachi Command Suite Replication Manager User Guide.*

Adding command devices

A command device supports replication operations by providing a communication channel between hosts and storage for replication commands.

Procedure

1. From the **Actions** menu, select **Manage Replication** to start the Replication Manager.

2. From the **Explorer** menu, select **Resources** and then **Storage Systems**.
3. Expand the tree and select the desired storage system.
4. Click the **Open** link, and then on the **Cmd Devs** tab, click **Add Cmd Devices** and add the command device.

Result

The added command device can be viewed by clicking the Open link from the storage system name shown in the Replication Manager. The list of added command devices are displayed in the Cmd Devs tab.

Related tasks

- Defining copy pairs on page 282
- Changing the status of a copy pair on page 283
- Deleting command devices on page 285

Editing command devices

You can edit an existing command device.

Procedure

1. From the **Actions** menu, select **Manage Replication** to start Replication Manager.
2. From the **Explorer** menu select **Resources** and then **Storage Systems**.
3. Expand the tree and select the desired storage system.
4. Click the **Open** link and then on the **Cmd Devs** tab click the pencil and paper icon and edit the command device.
5. Follow the instructions in the window to edit the command device.

Result

The edited command device can be viewed by clicking the Open link from the storage system name in Replication Manager. The list of edited command devices are displayed in the Cmd Devs tab.

Related tasks

- Defining copy pairs on page 282
- Changing the status of a copy pair on page 283

Related references

- HCS copy pair management operations (no Replication Manager license) on page 281

Deleting command devices

If you decommission the pair management server, delete the command device. When you delete command devices, the communication channel between hosts and storage for replication commands is deleted.

Procedure

1. From the **Actions** menu, select **Manage Replication** to start the Replication Manager.
2. From the **Explorer** menu, select **Resources** and then **Storage Systems**.
3. Expand the tree and select the desired storage system.
4. Click the **Open** link, and then on the **Cmd Devs** tab select command devices and click **Delete Cmd Devices**.

Result

The command devices you deleted no longer appear in the list of command devices.

6

Managing replication configurations using the Replication tab

This chapter describes how to replicate volumes for data protection (within or between storage systems) by using the Replication tab.

- ☐ About replicating volumes
- ☐ HCS Replication tab operations (with Replication Manager license)
- ☐ Managing replication configurations
- ☐ Managing copy topologies
- ☐ Monitoring replication configurations
- ☐ Replication and disaster recovery

About replicating volumes

Volume replication is used to make copies of critical data for volumes allocated to hosts or file servers/NAS modules. You can make a local copy of a volume in the same storage system, or make a remote copy of a volume in a different storage system.

In addition to Replication Manager, HCS includes the Replication tab. The Replication tab supports a convenient set of remote copy configurations, allowing the storage administrator to perform operations that would otherwise require accessing the Replication Manager GUI. (HRpM license required.)

Related concepts

- Allocating storage on page 141
- About analyzing Universal Replicator performance on page 552

Related tasks

- Creating a replication environment on page 293
- Discontinuing a replication configuration on page 300

Related references

- HCS Replication tab operations (with Replication Manager license) on page 288

HCS Replication tab operations (with Replication Manager license)

The Replication tab incorporates key functions that would otherwise require use of the Replication Manager GUI, including the creation and management of remote copies.

Note: Only open volumes are supported.

Operations supported

The Replication tab currently supports the following operations:

- **Setting up a replication configuration**
  From a single window, you can configure the necessary resources:
  - Remote paths
  - Quorum disks (GAD only)
  - Virtual storage machines (GAD only)
  - DP pools

- Journal groups (UR only)
- Thin Image pools
- Pair management servers

- **Refreshing configuration information**
  You can refresh the configuration information maintained by Replication Manager (without invoking Replication Manager).
- **Creating a copy topology**
  Using the Create Copy Topology wizard, you can create a remote copy using existing primary volumes. For secondary volumes, you can create new ones or you can select existing volumes and assign them to the host. The following configurations are supported:
  - 2DC Remote Copy
  - 3DC Cascade
  - 3DC Multi Target
  - 3DC Multi Target with Delta Resync
  - 2DC GAD
  - 3DC GAD+UR with Delta Resync

  Note: When you create a copy pair by using volumes with the capacity saving function enabled (dedupe and compression), copy or host I/O performance might be lowered. For details, see the manual for the applicable volume replication functionality.

- **Expanding or shrinking a copy topology**
  You can perform two types of expanding and shrinking operations:
  - *Add/delete copy groups or copy pairs*: you can add or delete pairs within an existing copy topology.
  - *Change the copy topology type*: you can expand or shrink a copy topology by altering the type.

  You can transition from one copy topology type to another by using the Expand Copy Topology wizard and Shrink Copy Topology window. (See About changing the copy topology type on page 307 for a list of supported transitions.)
- **Configuring alerts**
  You can configure alerts for copy groups, journal groups, and TI pools (without invoking Replication Manager).
- **Checking the status of a replication**
  You can detect problems by periodically checking the replication status and alerts.
- **Changing the status of a replication**
  You can change the status of a replication as necessary when a performing a takeover or takeback of a remote copy (backup or restore of a local copy required), or when you want to test disaster recovery.
- **Creating a GAD copy pair**

If a single instance of Device Manager manages all the storage systems that make up the GAD, refer to the procedures in Replicating volumes using global-active device on page 333 rather than this chapter. (Refer to Supported storage system and copy type combinations on page 291.)

Note: To back up file systems created on a file server or NAS Module, use the appropriate management software.

Operations not currently supported

Certain copy pair operations are not currently supported in the Replication tab workflow. Where possible, the procedures in this module include instructions for when to use the Replication Manager or Device Manager GUI in conjunction with the Replication tab.

- *Creating a cascade configuration for a local copy.* Use Replication Manager or Device Manager - Storage Navigator instead.
- *Configuring a local copy beyond eight generations.* Use Replication Manager instead.
- *Creating a local copy only (without a remote copy).* Use Replication Manager instead.
- *Creating P-VOLs for a remote site.* Instead, use the Resource tab of the remote Device Manager to select a storage system or host to allocate P-VOLs.
- *Deleting P-VOLs.* Instead, use the Resource tab to select a storage system or host to unallocate P-VOLs.
- Expanding to a 2DC remote copy (UR) configuration, and shrinking to a 2DC remote copy (UR) configuration. Use Replication Manager instead.
- Shrinking from a GAD+UR with Delta Resync configuration to a global-active device configuration. Use Replication Manager instead.
- Changing the copy pair status of an unsupported copy topology. Use the Change Pair Status Wizard from Replication Manager instead. (Refer to the table below for a list of supported copy types.)
- *Adding new copy pairs within an existing copy topology that yields an inconsistent number of copy pairs between copy groups.* Use Replication Manager instead. (Adding copy pairs to existing groups is supported.)
- *Deleting part of a copy topology that yields an inconsistent number of copy pairs between copy groups.* Use Replication Manager to delete part of a copy pair in the copy topology. (Deleting all related pairs and S-VOLs is supported.)
- Setting alerts for copy pair status and copy license usage. Use Replication Manager instead.

Supported storage system and copy type combinations

Primary storage system	Secondary storage system					
	Remote copy					Local copy
	VSP G1000, VSP G1500, VSP F1500	VSP Gx00 models[3]	VSP Fx00 models[4]	Virtual Storage Platform	HUS VM	Same storage system as primary storage system[1]
VSP G1000, VSP G1500, VSP F1500	TC, UR, GAD[2]	TC, UR	TC, UR	TC, UR	TC, UR	SI, TI
VSP Gx00 models[3],	TC, UR	TC, UR, GAD[2]	TC, UR	-	TC, UR	SI, TI
VSP Fx00 models[4]	TC, UR	TC, UR	TC, UR	-	TC, UR	SI, TI
Virtual Storage Platform	TC, UR	-	-	TC, UR	TC, UR	SI, TI
HUS VM	TC, UR	TC, UR -	TC, UR	TC, UR	TC, UR	SI, TI

Legend:

GAD: global-active device

SI: ShadowImage

TI: Thin Image

TC: TrueCopy

UR: Universal Replicator

Notes:

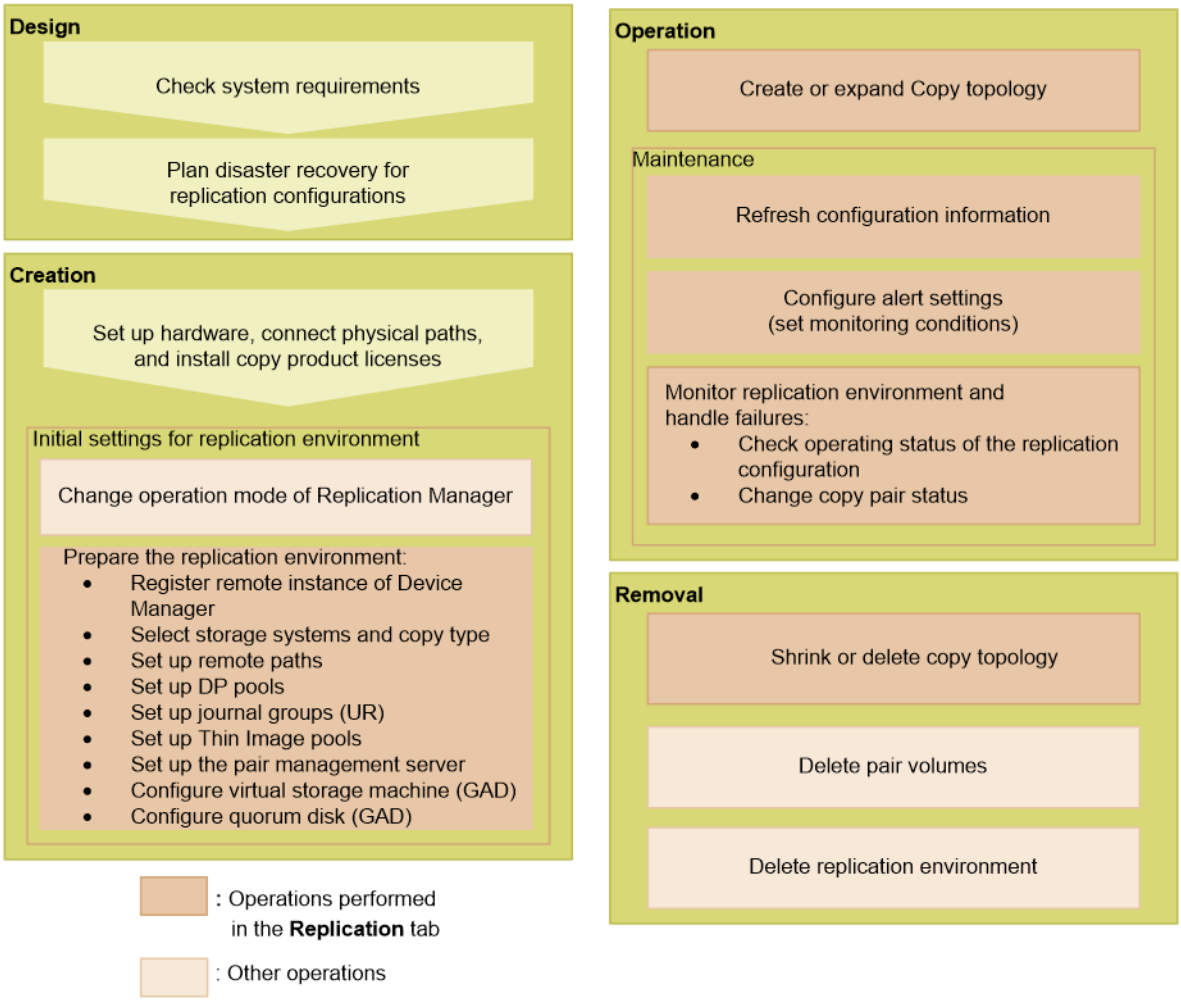
1. TI managed by snapshot groups is supported. For TI, copy pairs with the clone attribute and copy pairs that support cascade configurations (copy pairs for which the secondary volume is a DP volume) are not supported. To create or delete such a copy pair, use Device Manager - Storage Navigator.
2. For the secondary storage system of a Universal Replicator instance in a global-active device 3DC delta re-sync configuration, only VSP G1000, VSP G1500, VSP F1500, VSP G800, or VSP F800 can be used.
3. 3DC configuration only supported for VSP G800
4. 3DC configuration only supported for VSP F800

Managing replication configurations

The Replication tab can be used to create or discontinue a replication environment.

Volume replication configuration workflow

The following workflow describes the necessary planning and implementation steps for creating, monitoring, and managing a replication configuration.



For detailed information on site requirements, hardware setup, operational detail, and recovery recommendations, see the following:

- *Hitachi Command Suite Administrator Guide*
- *Hitachi Command Suite System Requirements*

Creating a replication environment

Use the Set up Replication/GAD window of the Replication tab to set up a new replication environment, including:

- Remote paths
- Quorum disks (GAD only)
- Virtual storage machines (GAD only)
- DP pools
- Journal groups (UR only)
- Thin Image pools
- Pair management servers

Before you begin

The local and remote storage systems have been added to HCS.

Procedure

1. In the **Replication** tab, **General Tasks** select **Set up Replication/GAD**.
2. (Optional) Click **Add Device Manager** to add an instance of Device Manager (and select a storage system managed by that instance).
3. Click **Select Storage Systems and Copy Type** to select the primary and secondary storage systems for which to build a replication configuration, and the type of copy to be used between the storage systems.

   If the copy type is global-active device, you should also select the virtual storage machine.
4. Click **OK**.
5. Perform the steps in each of the **Set up Replication/GAD** panels as necessary (see **Related tasks**).
6. Confirm that "Complete" is displayed as the status for all steps performed.

   Note: Even if "Complete" is displayed, you should verify the status of each resource before attempting to create a copy topology.
7. Click **Close**.
8. Follow the procedure in Creating a copy topology for a replication configuration (TC/UR) on page 303.

Result

You can check the execution result of each stage from the Set up Replication/GAD window.

Related tasks

- Creating remote paths for a replication configuration on page 294
- Configuring a quorum disk for a replication configuration on page 298
- Configuring a virtual storage machine for a replication configuration on page 299
- Configuring journal groups for a replication configuration on page 296
- Creating DP pools for a replication configuration on page 295
- Creating Thin Image pools for a replication configuration on page 297
- Configuring pair management servers for a replication configuration on page 297

Creating remote paths for a replication configuration

Use the Configure Remote Paths panel to create the necessary paths between the primary and secondary storage systems.

Before you begin

The primary and secondary storage systems have been selected.

Note: For more information about remote paths, see the *Hitachi Command Suite Replication Manager User Guide*.

Procedure

1. On the **Replication tab**, **General Tasks**, select **Set up Replication/ GAD**.
2. From the **Configure Remote Paths** panel, click **Create Remote Path**. The **Create Remote Path** wizard opens.
3. Check the copy type, primary storage system, and secondary storage system displayed.
4. (Optional) If TC is selected as the copy type, select the **Specify MCU/RCU check box** if necessary.
5. Click **Next**.
6. (Optional) Enter the labels in the **Primary -> Secondary** area.
7. From the list of path group IDs, select the ID of the path group to be assigned to the remote paths.
8. From the list of ports, select the ports to be used as the primary and secondary port.
9. If UR is selected as the copy type or TC is selected to create a bidirectional path, fill in the **Primary <- Secondary** area in the same way as in the **Primary -> Secondary** area. (If TC is selected and a bidirectional path is not created, clear the check box.)
10. Click **Next**.
11. Check the settings on the Confirm page, select the check box for confirmation, and click **Confirm**.

12. Click **Finish**.
13. Click **Refresh Display**.
14. Confirm that the paths are included in the remote path listing table and that the panel shows that the status is Complete.

Result

The remote paths between the primary and secondary storage systems are configured.

If setup of the remote path failed, in General Tasks, click Replication Manager. In the **Delete Remote Paths** dialog box, make sure the Do not change port attributes check box is selected, and then delete the remote path.

Creating DP pools for a replication configuration

The Configure DP Pools panel creates the necessary pools.

Tip: If there is an entry with the value `Unknown` displayed in the DP pools list, this means that the pair management server might not be set, or that communication between Device Manager and the Device Manager agent has failed.

Before you begin

- The primary and secondary storage systems have been selected.
- Basic volumes have been created.

Procedure

1. On the **Replication tab**, **General Tasks**, select **Set up Replication/ GAD**.
2. Open the **Configure DP Pools** panel.
3. Click **Create Pool** for the primary storage system.
4. Select the Pool Type.
5. Select the parity group to be added to the list of parity groups, select **Add Parity Groups**, and then perform any necessary operations.
6. Click **Show Plan**.
7. Check the plan summary (and plan details, if necessary) and click **Submit**.
8. Repeat steps 3-7 to create the desired number of DP pools for the primary storage system.
9. Click **Create Pool** for the secondary storage system and repeat the same procedure as necessary. (If you manage the storage system from a remote Device Manager, perform this step by clicking **Actions** and then **Create Pool** on the remote instance.)
10. Select the **Tasks & Alerts** tab. (If the DP pool was created at the remote site, perform this step on the remote Device Manager.)

11. Confirm that the DP pool creation task has finished.
12. If the DP pool was created at the remote site, return to the **Configure DP Pools** panel and click **Refresh Configuration**.
13. Click **Submit**. (Note that the refresh may take some time to complete.)
14. Click **Refresh Display**.
15. Confirm that the created DP pools are included in the DP pool list table.

Configuring journal groups for a replication configuration

The Configure Journal Groups panel of Set up Replication/GAD creates journal groups.

Tip: If there is an entry with the value `Unknown` displayed in the journal group list, this means that the pair management server might not be set, or that communication between Device Manager and the Device Manager agent has failed.

Before you begin

- The primary and secondary storage systems have been selected.
- For VSP or HUS VM, basic volumes have been created.
- For VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, or VSP F1500, DP pools have been created.
- Identify the number of volumes to create.
- Identify volume capacities (for creating a journal group using a newly-created DP volume).

Procedure

1. On the **Replication tab**, **General Tasks**, select **Set up Replication/ GAD**.
2. From the **Configure Journal Groups** panel, click **Create Volumes** to create DP volumes for each storage system (if needed).
3. Click **Create Journal Group** for each storage system. Supply the information as prompted by the Create Journal Group Wizard and confirm your choices as instructed.
4. Click **Finish** to return to the **Configure Journal Groups** panel.
5. Click **Refresh Display**.
6. (Optional) Click **Create Alert Setting** to set up alerts with the Create Alert Setting Wizard.
7. If you created alerts, click the second instance of **Refresh Display** to update the table information.

Related tasks

- Setting alerts for a replication configuration on page 314

Creating Thin Image pools for a replication configuration

The Configure Thin Image Pools panel creates the necessary Thin Image pools.

Tip: If there is an entry with the value `Unknown` displayed in the Thin Image pools list, this means that the pair management server might not be set, or that communication between Device Manager and the Device Manager agent has failed.

Before you begin

- The primary and secondary storage systems have been selected.
- Basic volumes have been created.

Procedure

1. On the **Replication tab**, **General Tasks**, select **Set up Replication/ GAD**.
2. Open the **Configure Thin Image Pools** panel.
3. Click **Create Pool** for each storage system.
   Supply the information as prompted by the Create Pool Wizard and confirm your choices as instructed.
4. Click **Finish** to return to the **Configure Thin Image Pools** panel.
5. Click **Refresh Display**.
6. (Optional) Click **Create Alert Setting** to set up alerts with the Create Alert Setting Wizard.
7. If you created alerts, click the second instance of **Refresh Display** to update the table information.

Related tasks

- Setting alerts for a replication configuration on page 314

Configuring pair management servers for a replication configuration

The Configure Pair Management Servers panel includes options to download the Device Manager Agent, create the necessary command devices, and edit the pair management server settings.

Before you begin

- The primary and secondary storage systems are selected.
- The hosts to be used as pair management servers are connected to the storage systems.
- A DP pool has been created (if you are creating a new volume to be used as a command device).
- A volume has been created (if you are designating an existing volume to be used as a command device).

Procedure

1. On the **Replication tab**, **General Tasks**, select **Set up Replication/ GAD**.
2. From the **Configure Pair Management Servers** panel, make sure that CCI is installed on the hosts.
3. If necessary, click **Download** to obtain the Device Manager Agent for your platform and install the software. Make sure that the pair management servers are registered in Device Manager as hosts.

   If you are using a virtual machine as a pair management server and the virtualization server uses NPIV (N-Port ID Virtualization) HBAs, you must also register the virtualization server in HCS.
4. If a pair management server is registered on a remote Device Manager, click **Refresh Configuration**.
5. Click **Allocate Command Device** for the primary and secondary storage systems.
6. Select to **Create a New Volume** and supply the necessary information or **Use an Unallocated Volume** to select an available volume.
7. (Optional) Supply the LUN Path Settings.
8. Click **Submit**.
9. When your command device settings are complete, click **Refresh Display**.
10. (Optional) Click **Edit Snapshot Groups Setting** to manage snapshot groups in each storage system.
11. (Optional) When your snapshot settings are complete, click **Refresh Display**.

Configuring a quorum disk for a replication configuration

The Configure Quorum Disks panel configures an external volume as a quorum disk for a global-active device.

Before you begin

- The primary and secondary storage systems are selected.
- The quorum storage system is connected to the primary and secondary storage system by Fibre Channel or Fibre Channel over Ethernet (FCoE).
- The port attribute settings for the primary and secondary storage system.
- The external mapping volumes are already allocated to the external ports (when using an unmanaged external storage system).
- The LDEV ID has been identified (when using an existing external volume).

Procedure

1. If you are using a volume on a storage system **not** managed by the local Device Manager:

   a. Log in to the element manager of the external storage system.

b. Identify the volume to use as the quorum disk. (Make a note of the LUN.)
c. Log off from the element manager.

2. On the **Replication tab**, **General Tasks**, select **Set up Replication/ GAD**.
3. From the **Configure Quorum Disks** panel, you have these options:
   - If the volume is on a storage system managed by the local Device Manager, select **Use a managed external storage system** and continue with step 5.
   - If the volume is on a storage system **not** managed by the local Device Manager, select **Use an unmanaged external storage system** and continue with step 9.
   - Otherwise, select **Use an existing external volume**.
4. From both the primary and secondary storage systems, select an existing external volume that is mapped for quorum disk use.

   If you manage the storage system from a remote Device Manager, right-click the target storage system in the **Resources** tab of the management-side Device manager, add the quorum disk from **Remote Connections**, and continue with step 12.
5. Click **Select Volume** to choose the volume to be used for the quorum disk.
6. In the **Volume Selection** dialog box, select a volume.
7. (Optional) Click **Create Volume** if you wish to create a new volume.
8. Click **OK**.
9. Click **Virtualize Volume**. If the storage system is managed by a remote Device Manager, log in and virtualize the volume.
10. Click **Configure Quorum Disks** to assign a quorum disk ID to virtualized volume. (You must assign the same quorum disk ID to both the primary storage system and the secondary storage system.)
11. Click **Execute**.
12. (Optional) If any storage systems are managed remotely, click **Refresh Configuration** and then click **Execute** in the **Refresh Configuration** dialog box.
13. In the **Configure Quorum Disks** panel, click **Refresh Display**.
14. Review the table of quorum disks to confirm that the disk has been created.

Configuring a virtual storage machine for a replication configuration

The Configure Virtual Storage Machines panel configures a virtual storage machine for a global-active device.

Before you begin

- The primary and secondary storage systems are selected.
- The virtual storage machine exists in the primary or secondary storage system.

- Identify the primary and secondary storage system resources to be added to the virtual storage machine.

Note: For the host group or iSCSI target number, use a host group or iSCSI target number that has not been allocated to a storage port.

Procedure

1. On the **Replication tab**, **General Tasks**, select **Set up Replication/ GAD**.
2. From the **Configure Virtual Storage Machines** panel, if there are existing (matching) virtual storage machines on the primary and secondary storage systems, you can add resources to each virtual storage machine by clicking **Edit Virtual Storage Machine**.

   Otherwise, click **Create Virtual Storage Machine** and add a new virtual storage machine on the other storage system along with the necessary resources. (Be sure and use the same model name and serial number as you selected in the **Select Storage Systems and Copy Type** dialog.)
3. If you want to add resources to a virtual storage machine on a storage system managed by the remote Device Manager:
   a. Log in to the management-side Device Manager of the primary storage system.
   b. Click **Create Virtual Storage Machine** or **Edit Virtual Storage Machine** in the **Management** tab.
   c. Add the resources necessary for the virtual storage machine. Do the same on the secondary storage system using the same model name and serial number as the virtual storage machine you selected in the **Select Storage Systems and Copy Type** dialog.
   d. Return to the **Configure Virtual Storage Machines** panel.
4. If you created or edited a virtual storage machine at a remote site, click **Refresh Configuration** and then click **Submit** in the **Refresh Configuration** dialog.
5. Click **Refresh Display**.
6. Review the table of virtual storage machines to confirm that the virtual storage machine has been created.

Discontinuing a replication configuration

To discontinue the use of a replication configuration, you must remove the pair management server settings, journal groups, pools, remote paths, and the associated quorum disks and virtual storage machines (GAD only).

Before you begin

- Data on affected pairs has been saved.

- All copy pairs and copy groups for remote copy operations between the two storage systems (including the storage system to be removed) have been deleted.
- Before you delete DP pools, first unallocate the constituent DP volumes.

Procedure

1. Delete the pair management server settings.
   a. Log in to Device Manager on the local or remote management server.
   b. Select the **Resources** tab.
   c. Select **Open-Allocated** for the storage system to be removed.
   d. Filter the volume table by using the following conditions:
      - the command device attribute is assigned
      - the host to which the volumes are allocated is a pair management server
   e. Select the filtered volumes and then click **Unallocate Volumes**.
   f. Check the details displayed in the **Unallocate Volumes** dialog box.
      Note the name of LDEV ID that will be unallocated.
   g. Execute the Unallocate Volumes task.
   h. (Optional) Start Replication Manager.
   i. (Optional) In the Storage Systems view, select the storage system to be removed.
   j. (Optional) In the **Cmd Devs** tab, select the LDEV ID that has been unallocated and click **Delete Cmd Devices**.
   k. (Optional) Delete the volumes.
   l. (Optional) If the server is no longer to be used as a pair management server, uninstall Device Manager Agent and CCI.
2. Delete the Thin Image pools.
   a. Start Replication Manager.
   b. In the Storage Systems view, select the storage system to be removed.
   c. Open the **Pools** tab and sort the table by the "Used by" column.
   d. Select the Thin Image pools and click **Delete Pools**.
   e. Check the details displayed in the **Delete Pools** dialog box and execute the operation.
3. Delete the journal groups.
   a. In the Storage Systems view, select the storage system to be removed.
   b. Open the **JNLGs** tab and sort the table by the "Used by" column.
   c. Select the journal groups and click **Delete JNLGs**.
   d. Check the details displayed in the **Delete JNLGs** dialog box and execute the operation.
4. Delete the DP pools.
   a. Log in to Device Manager on the local or remote management server.
   b. Select the **Resources** tab.
   c. Select the **DP pools** of the storage system to be removed.

d. Select the DP pools and click **Delete Pools**.
e. Check the details displayed in the **Delete Pools** dialog box and execute the operation.

5. Delete the remote paths.
   a. Start Replication Manager.
   b. In the Storage Systems view, select the storage system to be removed.
   c. Open the **Remote Paths** tab.
   d. Select the path group and click **Delete Paths**.
   e. Check the details displayed in the **Delete Remote Paths** dialog box and execute the operation.
6. For a GAD configuration, delete the quorum disk.
   a. In the **Resources** tab, select **Storage Systems**.
   b. Expand the tree and right click the target storage system. Select the displayed **Remote Connections**.
   c. In the **Quorum disk** tab, select the Quorum disk whose settings you want to delete, confirm it is correct, and delete it.
   d. From the **Actions** menu, select **Unvirtualize Volumes** and then, from the external storage system used for the Quorum disk, delete the virtualization of the volume (Quorum disk) for the target storage system.
   e. Check the plan and execute it. The plan is registered as a task.
   f. In the **Tasks & Alerts** tab, check the execution results of the task.
7. For a GAD configuration, delete the virtual storage machine
   a. In the **Administration** tab, select **Virtual Storage Machine**.
   b. From the list of virtual storage machines, select the target virtual storage machine and then click the **Edit Virtual Storage Machine** button.
   c. In the **Storage Systems** tab, select the storage system you want to delete from the virtual storage machine, and then click the **Remove Storage Systems** button.
   d. Check the plan and execute it. The plan is registered as a task.
   e. In the **Tasks & Alerts** tab, check the execution results of the task.

Result

The replication environment has been removed.

Related tasks

- Creating a replication environment on page 293

Site management with the Replication tab

By default, the Replication tab creates a list of sites based on the storage systems that are managed by Device Manager (and discovered by Replication Manager). If you have a large number of storage systems, you can create new sites (or edit the ones created automatically) to organize your storage systems (and their copy topologies) into logical groups.

You can manage sites from the Sites table of the Replication tab:

- Create Site
- Edit Site
- Delete Site

Tip: If you want to delete a site, you must first move the assigned storage systems to another site.

Managing copy topologies

The Replication tab can be used to create, expand, or shrink the copy topology of a replication configuration.

Creating a copy topology for a replication configuration (TC/UR)

You can use the Create Copy Topology wizard of the Replication tab to set up a simple 2DC configuration or a 3DC configuration (cascade, multi-target, or delta resync) using the TrueCopy or Universal Replicator copy types.

Before you begin

- Set up the replication environment.
- Update the information of the target primary and secondary storage systems.
- Allocate the primary volume to the host.
- Allocate the secondary volume to the host (when specifying an existing secondary volume).

Tip: While selecting the secondary volume:

- If you create a new host group or iSCSI target, the default host mode and host mode options are set to match the primary side. (To specify different values, clear the Apply the same host mode and host option as primary side check box.) Select the target volume in the Resource tab as necessary, and then edit the LUN path.
- If there is an entry with the value `Unknown` displayed for Used Capacity in the **Select Volume** window, this means that the pair management server might not be set, or that communication between Device Manager and the Device Manager agent has failed.

Procedure

1. In the **Replication** tab, click **Create Copy Topology**.
2. On the **Introduction** page, click **Next**.
3. On the **Select Copy Topology Type** page, select a (non-GAD) copy topology and click **Next**.

4. On the **Configure Copy Topology** page, click the **Site Storage System** icon for the Primary site.
   Select a storage system and click **OK**.
5. Click the **Volumes** icon for the Primary site.
   Select the volumes to be used as P-VOLs and click **OK**.
6. If you are creating a 2DC configuration, skip to step 10.
7. Click the **Site Storage System** icon for the Intermediate site.
   Select a storage system and click **OK**.
8. Click the **Volumes** icon for the Intermediate site.
   Create new volumes by selecting the DP pool and the LUN path (Storage Port), or select existing volumes to be used as S-VOLs, and click **OK**.
9. Click the **Copy Group** icon between the Primary and Intermediate sites.
   Provide the necessary information and click **OK**.
10. Click the **Site Storage System** icon for the Secondary site.
    Select a storage system and click **OK**.
11. Click the **Volumes** icon for the Secondary site.
    Create new volumes by selecting the DP pool and the LUN path (Storage Port), or select existing volumes to be used as S-VOLs, and click **OK**.
12. If you are creating a Cascade configuration, skip to step 16.
13. Select the **Copy Group** icon between the Primary and Secondary sites.
    Provide the necessary information and click **OK**.
14. Click **Next**.
15. If you are creating a 2DC configuration, skip to step 17.
16. Select the **Copy Group** icon between the Intermediate and Secondary sites.
    Provide the necessary information and click **OK**.
17. If you want to add a local copy for performing a backup or disaster recovery test, click the **Add Local Copies** icon for each site, set up the copy groups and secondary volumes, and click **OK**.
18. Click **Next**.
19. (Optional) Change the execution **Schedule**.
20. Review the contents of the **Show Plan** page and click **Submit**.

Result

You can check the execution result from the Tasks & Alerts tab.

Related tasks

- Expanding a copy topology by adding a copy pair on page 310
- Shrinking a copy topology by deleting a copy pair on page 312
- Adding and deleting copy pairs in an unsupported copy topology on page 313

Creating a copy topology for a replication configuration (GAD)

You can use the Create Copy Topology wizard of the Replication tab to set up a standard 2DC GAD configuration or a 3DC GAD+UR with delta resync using global-active device.

Before you begin

- Set up the replication environment.
- Update the information of the target primary and secondary storage systems.
- Allocate the primary volume to the host.
- Allocate the secondary volume to the host (when specifying an existing secondary volume).

Tip: While selecting the secondary volume:

- If the host operating system is Windows or AIX and you are using the MPIO function, set the host mode option of the host group to 102. (You can do this by editing the LUN path in the Resources tab.)
- If you create a new host group or iSCSI target, the default host mode and host mode options are set to match the primary side. (To specify different values, clear the Apply the same host mode and host option as primary side check box.) Select the target volume in the Resource tab as necessary, and then edit the LUN path.
- If there is an entry with the value `Unknown` displayed for Used Capacity in the **Select Volume** window, this means that the pair management server might not be set, or that communication between Device Manager and the Device Manager agent has failed.

Procedure

1. Open the **Replication** tab.
2. (Optional) In **General Tasks**, select the **Allocate Volumes** button to create a new volume or assign an existing volume to be used as a P-VOL in the copy topology that you want to configure (if the primary storage system is added to Management Server (Local) and you need to create or allocate a new P-VOL).
3. In **General Tasks**, select **Create Copy Topology**.
4. On the **Introduction** page, click **Next**.
5. On the **Select Copy Topology Type** page, select the desired global-active device copy topology and click **Next**.
6. On the **Configure Copy Topology** page, click the **Site Storage System Virtual Storage Machine** icon for the Primary site.
7. From the **Storage System** drop-down list, select the storage system to be used as the primary storage system to narrow down the virtual storage machines.

8. From the **Virtual Storage Machine** drop-down list, select a virtual storage machine.
9. Click **OK**.
10. Click the **Volumes** icon for the Primary site.
    Select the volumes to be used as P-VOLs and click **OK**.
11. If you are creating a GAD 2DC configuration, skip to step 15.
12. Click the **Site Storage System Virtual Storage Machine** icon for the Intermediate site.
    Select a storage system and click **OK**.
13. Click the **Volumes** icon for the intermediate site.
14. Click the **Copy Group** icon between the Primary and Intermediate sites.
    Provide the necessary information and click **OK**.
15. Click the **Site Storage System** icon for the Secondary site.
    Select a storage system and click **OK**.
16. Click the **Volumes** icon for the Secondary site.
    Create new volumes by selecting the DP pool and the LUN path (Storage Port), or select existing volumes to be used as S-VOLs, and click **OK**.
17. If you are creating a 2DC configuration, skip to step 20.
18. Select the **Copy Group** icon between the Intermediate and Secondary sites.
    Provide the necessary information and click **OK**.
19. Click **Next**.
20. Select the **Copy Group** icon between the Primary and Secondary sites.
    Provide the necessary information and click **OK**.
21. If you want to add a local copy for performing a backup or disaster recovery test, click the **Add Local Copies** icon for each site, set up the copy groups and secondary volumes, and click **OK**.
22. Click **Next**.
23. (Optional) Change the execution **Schedule**.
24. Review the contents of the **Show Plan** page and click **Submit**.

Result

You can check the execution result from the Tasks & Alerts tab.

Related tasks

- Expanding a copy topology by adding a copy pair on page 310
- Shrinking a copy topology by deleting a copy pair on page 312
- Adding and deleting copy pairs in an unsupported copy topology on page 313

Changing the copy topology name

You can use the Replication tab to change the name of a copy topology.

Note: The following can affect the copy topology name:

- If you change a copy group name of a remote copy by using Replication Manager or CCI, the name of the copy topology in the Replication tab might change to incorporate the new the copy group name.
- The copy topology name might be returned to the default if you change the WWN/iSCSI name (for example, if you replace an HBA of the pair management server after changing the name of the copy topology).

If this occurs, you can change the name back.

Procedure

1. Select the **Replication** tab.
2. From the **Copy Topology** table, select the desired entry.
3. Select **Manage Copy Topologies > Edit Copy Topology Name**.
4. Change the name and click **Execute**.
5. Check the result in the **Edit Copy Topology Name Complete** window, and click **OK**.
6. Check the updated **Copy Topology** table to confirm the new name.

About changing the copy topology type

In addition to adding or removing copy pairs and copy groups in a copy topology, you can also expand or shrink a copy topology by changing the copy topology type.

The following tables list all the supported copy type transitions for expanding and shrinking copy topologies.

Table 15 Copy topology types (expansion)

Copy topology type prior to expansion	Copy topology type after expansion			
	3DC Cascade	3DC Multi Target	3DC Multi Target with Delta Resync	GAD+UR with Delta Resync
2DC Remote Copy (TrueCopy)	Y	Y	Y[1]	-
2DC Remote Copy (Universal Replicator)	N	N	N	-
3DC Cascade	-	-	Y[2]	-
3DC Multi Target	-	-	Y	-
global-active device	-	-	-	Y

Legend:

Y = Supported by Replication tab

N = Not supported (use Replication Manager)

- = Not applicable

Copy topology type prior to expansion	Copy topology type after expansion			
	3DC Cascade	3DC Multi Target	3DC Multi Target with Delta Resync	GAD+UR with Delta Resync

Notes:

1. First expand from a 2DC Remote Copy (TrueCopy) configuration to a 3DC Multi Target, then finally expand to a to 3DC Multi Target with Delta Resync configuration.
2. First shrink from a 3DC Cascade configuration to a 2DC Remote Copy (TrueCopy) , then expand to a 3DC Multi Target, and finally expand to a 3DC Multi Target with Delta Resync configuration.

Table 16 Copy topology types (shrinking)

Copy topology type prior to shrinking	Copy topology type after shrinking				
	2DC Remote Copy (TrueCopy)	2DC Remote Copy (Universal Replicator)	3DC Cascade	3DC Multi Target	global-active device
3DC Cascade	Y	N	-	-	-
3DC Multi Target	Y	N	-	-	-
3DC Multi Target with Delta Resync	Y[1]	N	Y[2]	Y	-
GAD+UR with Delta Resync	-	-	-	-	N

Legend:

Y = Supported by Replication tab

N = Not supported (use Replication Manager)

- = Not applicable

Notes:

1. First shrink from a 3DC Multi Target with Delta Resync configuration to a 3DC Multi Target, then finally shrink to a 2DC Remote Copy (TrueCopy) configuration.
2. First shrink from a 3DC Multi Target with Delta Resync configuration to a 3DC Multi Target, then shrink to a 2DC Remote Copy (TrueCopy), and finally expand to a 3DC Cascade configuration.

Expanding a copy topology by changing the type

You can use the Expand a Copy Topology wizard of the Replication tab to change the copy topology type and expand the capacity.

Before you begin

Refer to About changing the copy topology type on page 307 to determine which copy type transitions are supported by the Replication tab and note any special requirements. (For example, some transitions require a combination of expanding and shrinking to achieve the desired

configuration.) If the transition is not supported, you will have to use the Replication Manager instead.

Procedure

1. Open the **Replication** tab.
2. Select the existing entry from the **Copy Topology** table.
3. Select **Manage Copy Topologies > Expand a Copy Topology**.
4. On the **Select Copy Topology Type** page, choose the expanded copy topology you wish to implement.
5. On the **Configure Copy Topology** page, click the **Site Storage System** icon for the secondary site.

   If you are expanding from 3DC Multi Target to 3DC Multi Target with Delta Resync, skip to Step 7.
6. Click the **Volumes** icon for the secondary site.

   Create new volumes by selecting the DP pool and the LUN path (Storage Port), or select existing volumes to be used as S-VOLs, and click **OK**.
7. Click the **Copy Group** icons to set the copy groups according to your desired configuration.
8. Provide the necessary information and click **OK**.
9. If you want to add a local copy for performing a backup or disaster recovery test, click the **Add Local Copies** icon for each site, set up the copy groups and secondary volumes, and click **OK**.
10. Click **Next**.
11. (Optional) Change the execution **Schedule**.
12. Review the contents of the **Show Plan** page and click **Submit**.

Result

You can check the execution result from the Tasks & Alerts tab.

Shrinking a copy topology by changing the type

You can use the Shrink a Copy Topology function of the Replication tab to change the copy topology type and shrink the capacity.

Before you begin

Refer to About changing the copy topology type on page 307 to determine which copy type transitions are supported by the Replication tab and note any special requirements. (For example, some transitions require a combination of expanding and shrinking to achieve the desired configuration.) If the transition is not supported, you will have to use the Replication Manager instead.

Procedure

1. Open the **Replication** tab.

2. Select the existing copy topology from the **Copy Topology** table.
3. Select **Manage Copy Topologies > Shrink a Copy Topology**.
4. On the **Shrink a Copy Topology** page, select the revised topology (**Copy topology type (shrunk)**) that you wish to implement.
5. (Optional) You can choose to delete secondary volumes of local and remote copy groups.
6. Click **Show Plan**.
7. Review the contents of the **Show Plan** page. Verify the copy groups and copy pairs to be deleted.
8. Click **Submit**.

Result

You can check the execution result from the Tasks & Alerts tab.

Expanding a copy topology by adding a copy pair

You can use the Replication tab to add a copy pair to an existing copy topology without adding a copy group.

Before you begin

- Allocate the primary volume to the host.
- Create a DP pool when using a newly-created secondary volume (except for Thin Image).
- Allocate the secondary volume to the host (when specifying an existing secondary volume).

Tip: While selecting the secondary volume:

- For GAD configurations, if the host operating system is Windows or AIX and you are using the MPIO function, set the host mode option of the host group to 102. (You can do this by editing the LUN path in the Resources tab.)
- If you create a new host group or iSCSI target, the default host mode and host mode options are set to match the primary side. (To specify different values, clear the Apply the same host mode and host option as primary side check box.) Select the target volume in the Resource tab as necessary, and then edit the LUN path.
- If there is an entry with the value `Unknown` displayed for Used Capacity in the **Select Volume** window, this means that the pair management server might not be set, or that communication between Device Manager and the Device Manager agent has failed.

Procedure

1. Select the **Replication** tab.
2. From the **Copy Topology** table, select the entry to be expanded.
3. Select **Manage Copy Topologies > Add Copy Pairs**.

4. If the copy topology is not supported by the **Replication** tab, see Adding and deleting copy pairs in an unsupported copy topology on page 313.
5. Select the Primary **Volumes** and choose a volume to be added as the P-VOL.
6. Click **OK**.
7. Select each **Volume** box that lacks a green checkmark and add necessary volumes. (The **Next** button will not be available until all the boxes are checked. Warning symbols are displayed if there is a mismatch in your configuration.)
8. Click **Next**.
9. Check the contents of the **Show Plan** page.
10. (Optional) Change the execution **Schedule**.
11. Click **Submit**.

Related tasks

- Expanding a copy topology by adding a local copy group on page 311

Expanding a copy topology by adding a local copy group

You can use the Replication tab to add a local copy group to an existing copy topology.

Before you begin

- Set up the replication environment.
- Allocate the secondary volume to the host (when specifying an existing secondary volume).



Tip: While selecting the secondary volume:

- If you create a new host group or iSCSI target, the default host mode and host mode options are set to match the primary side. (To specify different values, clear the Apply the same host mode and host option as primary side check box.) Select the target volume in the Resource tab as necessary, and then edit the LUN path.
- If there is an entry with the value `Unknown` displayed for Used Capacity in the **Select Volume** window, this means that the pair management server might not be set, or that communication between Device Manager and the Device Manager agent has failed.

Procedure

1. Select the **Replication** tab.
2. From the **Copy Topology** table, select the entry to be expanded.
3. Select **Manage Copy Topologies > Add Copy Pairs**.
4. Select the **Add Local Copy** box for each site to which a local copy pair is to be added.
5. Add local copy groups (up to 8) and enter a name.

If you choose to create new volumes, be sure and enter (or select) a **Host Group** name. If you choose to use existing volumes, be sure and select the primary volume and click **Select secondary volume** to select the secondary.

6. Click **OK** when your copy group settings are complete.
7. Make sure that there are no warnings for any boxes in the copy topology.
8. Click **Next**.
9. Check the contents of the **Show Plan** page.
10. (Optional) Change the execution **Schedule**.
11. Click **Submit**.

Expanding an unsupported copy topology

If you want to expand a copy topology that is not supported by the Replication tab, you must use Replication Manager instead.

Shrinking a copy topology by deleting a copy pair

You can use the Replication tab to delete a set of copy pairs associated with a P-VOL in a copy topology without deleting a copy group.

Tip: If you specify that an S-VOL is to be deleted at the same time as the copy pair, the delete task might fail during the storage system refresh task. In such a case, use the Resources tab to first unallocate the volumes and then delete them.

Procedure

1. Select the **Replication** tab.
2. From the **Copy Topology** table, select the entry to be modified.
3. Select **Manage Copy Topologies > Delete Copy Pairs**.
4. If the copy topology is not supported by the **Replication** tab, a warning message is displayed and you must use Replication Manager instead as described in Adding and deleting copy pairs in an unsupported copy topology on page 313.
5. Select the volumes targeted for copy pair deletion from the list of **Top Primary Site** P-VOLs in the copy topology and click **Add**.
6. To delete the S-VOLs (remote or local) along with the copy pair, select the appropriate option at the bottom of the window.
7. Click **Show Plan**.
8. Check the target copy pairs and volumes for deletion shown in the plan confirmation window.
9. (Optional) Change the execution **Schedule**.
10. Click **Submit**.

Related tasks

- Shrinking a copy topology by deleting a local copy group on page 313

Shrinking a copy topology by deleting a local copy group

You can use the Replication tab to delete a copy pair or copy group of local copies in an existing copy topology.

Tip: If you specify that an S-VOL is to be deleted at the same time as the copy pair, the delete task might fail during the storage system refresh task. In such a case, use the Resources tab to first unallocate the volumes and then delete them.

Procedure

1. Select the **Replication** tab.
2. From the **Copy Topology** table, select the entry to be modified.
3. Select **Manage Copy Topologies > Delete Local Copy Groups**.
4. If the copy topology is not supported by the **Replication** tab, a warning message is displayed and you must use Replication Manager instead as described in Adding and deleting copy pairs in an unsupported copy topology on page 313.
5. Select the target copy group for deletion.
6. To delete the S-VOLs (remote or local) along with the copy pair, select the appropriate option at the bottom of the window.
7. Select **Show Plan**.
8. Check the target copy pairs and volumes for deletion shown in the plan confirmation window.
9. (Optional) Change the execution **Schedule**.
10. Click **Submit**.

Adding and deleting copy pairs in an unsupported copy topology

If you want to add or delete copy pairs in a copy topology that is not supported by the Replication tab, you must use Replication Manager as described here.

Procedure

1. In the **Copy Group** table of the **Replication** tab, check and make a note of the name of the copy group from which copy pairs are to be added or deleted and the pair management server that manages the copy group.
2. In **General Tasks**, click **Replication Manager**.
3. Based on the information from step 1, start the Pair Configuration Wizard for the relevant copy group.
4. Add or delete the copy pair(s).
5. Make sure that the Replication Manager task has finished.
6. Return to the **Replication** tab and refresh the display.

7. (Optional) Use the **Unallocate Volume** and **Delete Volume** dialogs to cancel volume allocation or remove the volume.

Deleting a copy topology from a replication configuration

You can delete a copy topology when is no longer necessary. You can delete the copy pairs included in a copy topology all at once.

Note: If you specify that an S-VOL is to be deleted at the same time as the copy pair, the delete task might fail during the storage system refresh task. In such a case, use the Resources tab to first unallocate the volumes and then delete them.

Procedure

1. Select the **Replication** tab.
2. From the **Copy Topology** table, select the entry to be deleted.
3. Select **Manage Copy Topologies > Delete Copy Pairs**.
4. Use Replication Manager to delete the copy topologies that cannot be deleted by using the **Replication** tab.
5. Select all the copy pairs, then click **Show Plan**.
6. Check the target copy pairs and volumes for deletion shown in the plan confirmation window.
7. (Optional) Change the execution **Schedule**.
8. Check the result of executing the task on the **Tasks & Alerts** tab.

Monitoring replication configurations

The Replication tab can be used to define alerts and monitor the status of a replication configuration.

Setting alerts for a replication configuration

You can set alerts for any copy group associated with a copy topology by selecting an entry in the Copy Groups table and clicking Create Alerts to start the Create Alert Setting Wizard.

You also can set up alerts for journal groups and Thin Image pools as you add them to a replication configuration.

Before you begin

If you intend to monitor the following resources using alerts, the resources must have been created and identified by Replication Manager, and displayed in the Replication tab:

- Copy groups
- Journal groups

- Thin Image pools

Tip: To modify existing alerts, go to the Alerts tab and click Edit Alert Setting. This opens the Alert Setting List in Replication Manager, where you can make the desired changes.

Procedure

1. When the Create Alert Setting Wizard opens (Introduction page), click **Next**.
2. On the **Select Monitoring Type** page:
   - if you want to set an alert for monitoring copy pair status, select **Pair Status Monitoring.**
   - if you want to set an alert for monitoring the usage of the journal volume capacity for each copy group, journal group, or the usage of Thin Image pools, select **Performance Monitoring.**
3. Click **Next**.
4. On the **Select Alert Setting** page, choose from the following:
   - **Create New Alert Setting**: to add a new alert. Specify the alert type, condition, and notification method as desired.
   - **Select Existing Alert Setting**: to add copy groups or copy pairs to an existing alert setting in the **Alert Setting List**.
   - **Alert Setting List**: to select an existing alert setting for which the target journal groups or Thin Image pools are to be added.
5. On the **Edit Alert Action** page, specify the alert type, condition, and notification method as desired.
   For **Performance Monitoring**, select the performance information to be used as an alert monitoring condition and specify the threshold for the monitored target.
6. Click **Next**.
7. Check the content of the Confirm page, select the check box for confirmation, and then click **Confirm**.
8. Click **Finish** to close the Create Alert Setting Wizard.
9. Check the updated **Copy Groups** list to confirm the alert settings.

Related tasks

- Monitoring the status of a replication configuration on page 315

Monitoring the status of a replication configuration

You can use the Replication tab to monitor alerts or check the status of a replication configuration and resolve any errors.

The main window of the Replication tab includes a summary of alerts. You can use the Alerts table to locate the affected resource, or check the status of copy pairs in the Copy Topology and Sites tables.

Before you begin

- Specify alert settings for the resources to be monitored.

Procedure

1. Select the **Replication** tab.
2. If you have received an alert, go the **Alerts** table and find the entry that matches the error received.

   If you are simply checking the status of a replication configuration, go to the **Copy Topology** table, or click the link for a site in the **Sites** table.
3. Click the link for the copy topology you wish to examine.
4. From the **Copy Groups** or **Copy Pairs** tables, locate and select the entry showing an error.
5. (Optional) Click **Change Pair Status** to correct the status of the copy group or copy pair.
6. (Optional) In the application area, click the **View HRpM Tasks** link to confirm the task has completed.

Related tasks

- Changing the status of a replication configuration on page 316
- Changing the status of an unsupported configuration on page 317

Changing the status of a replication configuration

You can use the Replication tab to change the status of copy pairs.

Procedure

1. Select the **Replication** tab.
2. Select an entry from the **Copy Topology** table.
3. Check the copy topology in the application area or the summary status of the copy groups in the **Copy Group** table.
4. In the **Copy Pair** table, check the copy pair status of the copy group with a problem.
5. Select the copy group or copy pair with a problem.
6. Click **Change Pair Status**.
7. If the copy topology is not supported by the **Replication** tab, see Changing the status of an unsupported configuration on page 317.
8. Correct the status of the copy group or copy pair.
9. In the application area, click the **View HRpM Tasks** link to confirm the task has completed.

Related tasks

- Changing the status of an unsupported configuration on page 317

Changing the status of an unsupported configuration

If the copy topology is not supported by the Replication tab, follow this procedure:

Procedure

1. Start Replication Manager.
2. Display the information about the copy group or container containing the copy pairs.
3. Click **Change Pair Status**.
4. Make sure that the Replication Manager task has finished.
5. Return to the **Replication** tab and refresh the display.

Replication and disaster recovery

The Replication tab can be used to perform disaster recovery and create local copies of a replication configuration for the purpose of testing.

Backing up and restoring a local copy

You can use the Replication tab to check the status of a replication configuration and change the pair status for the purpose of backing up or restoring a local copy.

For more information on backup and restore operations, see the *Hitachi Command Suite Replication Manager User Guide*.

Procedure

1. Select the **Replication** tab.
2. Click the link in the **Copy Topology** table.
3. In the **Copy Groups** table, select the copy group for which the operation is to be performed.
4. Click **Change Pair Status**.
5. To backup a local copy pair in a synchronized state, split the target copy pair or copy group.

   To backup a local copy pair that has already been split, resynchronize (in the forward direction) and then split (Take Snapshot) the target copy pair or copy group.

   To restore a local copy pair that has already been split, resynchronize (in the backward direction) the target copy pair or copy group.
6. In the application area, click the **View HRpM Tasks** link to confirm the task has completed.

Related tasks

- Performing replication takeover and takeback on page 318

Performing replication takeover and takeback

You can use the Replication tab to check the status of a replication configuration and change the pair status for the purpose of takeover or takeback of a remote copy.

For more information on takeover and takeback operations, see the *Hitachi Command Suite Replication Manager User Guide*.

Procedure

1. Select the **Replication** tab.
2. Click the link in the **Copy Topology** table.
3. Click **Change Pair Status**.
4. Execute one of the following advanced pair operations on the target copy group or copy pair: takeover, swap, takeover-recovery (recover with resync), or takeover-recovery (recover with recreate).
5. In the application area, click the **View HRpM Tasks** link to confirm the task has completed.
6. Make sure that the status of the copy pair, copy group, or copy topology is correct.

Related tasks

- Performing a disaster recovery test using a backup (split TC/UR copy) on page 318
- Performing a disaster recovery test using a snapshot (resync and split SI/TI copy) on page 319

Performing a disaster recovery test using a backup (split TC/UR copy)

You can use the Replication tab to change the pair status for the purpose of performing a disaster recovery test. In this scenario, a backup is taken (by splitting a TrueCopy/Universal Replicator pair).

Procedure

1. Select the **Replication** tab.
2. Click the link in the **Copy Topology** table.
3. Check the copy topology in the application area or the summary status of copy groups in the **Copy Group** table, and make sure that the status of the remote copy is Sync.
4. In the **Copy Topology** table, click the link for the copy topology.
5. Select the copy group of the remote copy to be used for the test.

6. Click **Launch Change Pair Status**.
7. Split the copy pairs in the copy group.
8. In the application area, click the **View HRpM Tasks** link to confirm the task has completed.
9. Make sure that the summary status of the copy group is Suspend.
10. Mount the S-VOLs of the copy group on the disaster recovery test server and conduct a test.
11. (Optional) Unmount the S-VOLs of the copy group from the disaster recovery test server.
12. Reselect the copy group in the **Copy Group** table.
13. Click **Change Pair Status**.
14. Resync the pairs of the copy group.
15. In the application area, click the **View HRpM Tasks** link to confirm the task has completed.
16. Make sure that the summary status of the copy group is Sync.

Performing a disaster recovery test using a snapshot (resync and split SI/TI copy)

You can use the Replication tab to change the pair status for the purpose of performing a disaster recovery test. In this scenario, a snapshot is taken (resync and split ShadowImage/Thin Image).

Procedure

1. Select the **Replication** tab.
2. Click the link in the **Copy Topology** table.
3. If a local copy is already available and an S-VOL has been allocated to the disaster recovery test server, skip to step 4. Otherwise, follow these steps:
   a. Click **Expand Copy Topology**.
   b. Click **Add Local Copy** and provide the required information.
   c. Click **OK**.
   d. Click **Next** and confirm the contents of the **Show Plan** page.
   e. (Optional) Change the execution **Schedule**.
   f. Click **Submit**.
4. In the **Copy Topology** table, click the link for the copy topology.
5. Select the copy group of the local copy associated with the remote copy.
6. Click **Change Pair Status**.
7. Take a snapshot (resync and split).
8. In the application area, click the **View HRpM Tasks** link to confirm the task has completed.
9. Make sure that the summary status of the copy group is Suspend.
10. Mount the S-VOLs of the copy group on the disaster recovery test server and conduct a test.

11. (Optional) Unmount the S-VOLs of the copy group from the disaster recovery test server.
12. (Optional) Delete the copy topology.

7

Managing virtual storage machines

Virtual storage machines are used to achieve continuous access to data for global-active device volumes and for data migration.

- ☐ About virtual storage machines
- ☐ Prerequisites for allocating resources to virtual storage machines
- ☐ Creating virtual storage machines
- ☐ Allocating volumes to hosts by using virtual storage machine resources
- ☐ Editing virtual storage machines
- ☐ Deleting virtual storage machines

About virtual storage machines

Global storage virtualization provides an abstraction layer between open systems hosts and storage systems. This layer enables non-disruptive storage management, such as:

- Non-disruptive host access to global-active device volumes which consist of a P-VOL and an S-VOL that reside on different storage systems.
- Non-disruptive host volume data migration from a source volume on an existing storage system to a target volume on a new storage system.

The features described above are accomplished through the use of virtual storage machines.

There are two types of virtual storage machines within the VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, and the VSP F1500 storage system:

- Default virtual storage machine
  One default virtual storage machine is created by default. The virtual information for the default virtual storage machine uses the same values as the physical information. Initially, all resources added to user-defined virtual storage machines are allocated from the default storage machine. Allocated resources are stored in the automatically created resource group (resource pool) when the virtual storage machine is created.
- User-defined virtual storage machines
  You can have a maximum of eight virtual storage machines for individual VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, or VSP F1500. This maximum includes the default virtual storage machine and seven user-defined virtual storage machines. Resources in the resource pools of user-defined virtual storage machines are added and removed by editing the virtual storage machines.

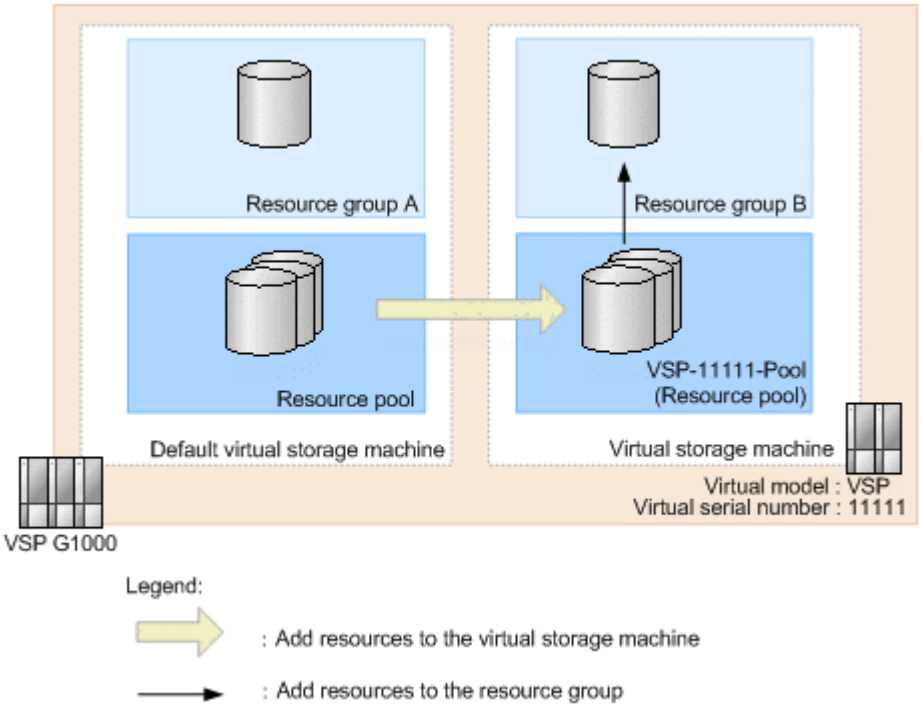
Virtual storage machine resources and resource groups

The following resources can be added and grouped by system resource type to the virtual storage machine:

- Parity groups
- LDEV IDs
- Storage ports
- Host group/iSCSI target numbers

User-defined resource groups defining more specific storage access can be created for the virtual storage machine. You create a resource group on a virtual storage machine by allocating resources from the resource pool. You can check the resource pools on separately created virtual storage machines from the resource group list.

The following figure shows an example configuration of a virtual storage machine.

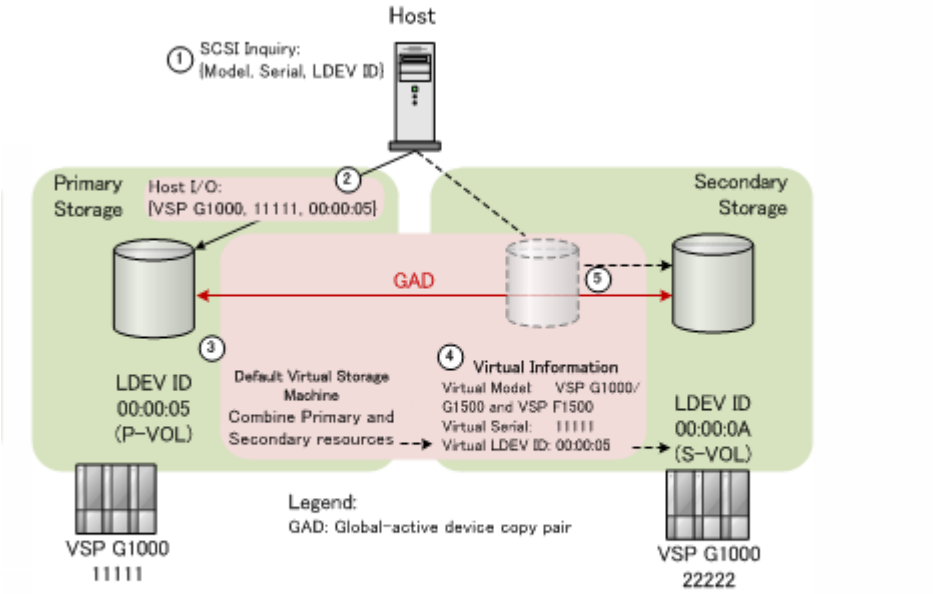


In the example configuration, the virtual storage machine is created using the model type and serial number of the VSP G1000. Resources from the resource pool of the default virtual storage machine are added to the user-defined virtual storage machine. The resources are added to the VSP-11111-Pool resource pool. When Resource Group B is created on the virtual storage machine, resources from the VSP-11111-Pool are added.

Identify the target resources by verifying the correlation between the virtual information that was sent from the host administrator and the physical information. You can name a virtual storage machine relative to how the virtual storage machine is being used.

Global-active device using virtual storage machines

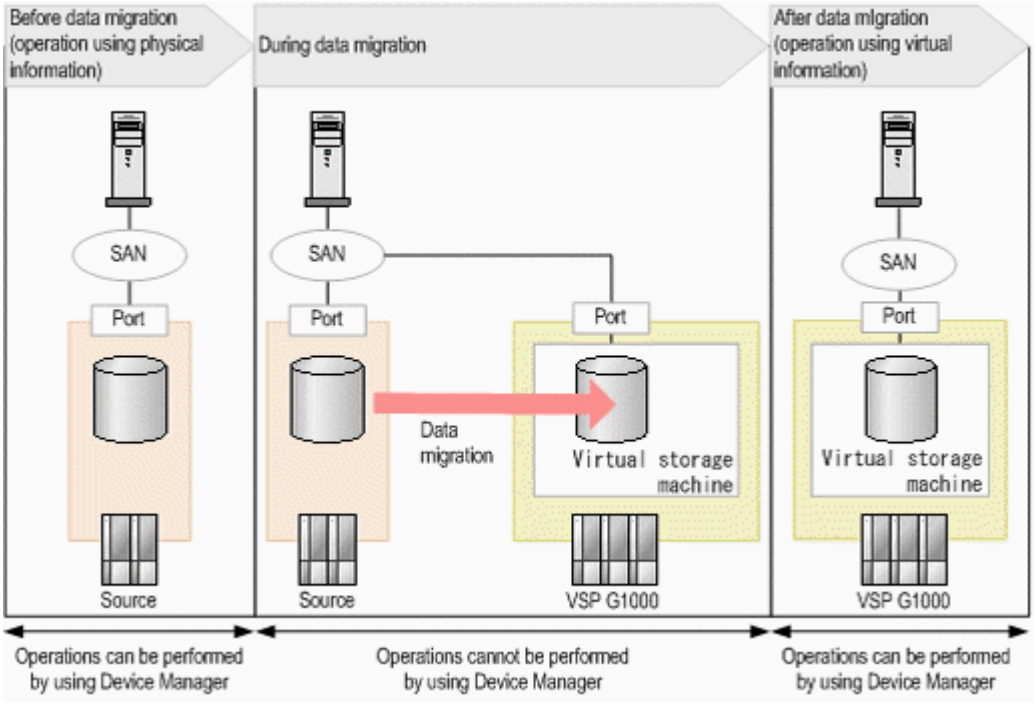
Global-active device employs volume replication (P-VOL and S-VOL in a bidirectional copy pair) to provide a high availability environment for hosts across storage systems and sites. The purpose of global-active device is to provide data protection and minimize data access disruptions for host applications due to storage system or site failures, as illustrated by the following figure.



1. Hosts use the SCSI Inquiry command to determine storage and device information for accessing volumes.
2. Host I/O is performed with this information.
3. During global-active device setup, the default virtual storage machine of the Primary storage is edited to include secondary storage system resources.
4. During volume allocation global-active device copy pairs are created and the virtual LDEV ID is mapped to the LDEV ID of the S-VOL. The virtual storage machine information is used to support data replication from P-VOL to S-VOL for the global-active device copy pair.
5. If the P-VOL becomes unavailable, the host accesses the S-VOL without disruption.

Data migration using virtual storage machines

The purpose of virtual storage machines when using nondisruptive migration functionality, is to provide continuous access to data during migration. The nondisruptive migration workflow includes creating one or more virtual storage machines on the target storage system. The virtual storage machine contains the information about the source storage system resources.



During the three stages of data migration, accessing the information changes from using physical to virtual information.

- Before data migration (using physical information)
  Before data migration, a host uses physical information to access the source storage system resources.
- During the data migration (using virtual information)
  Configure virtual storage machine information on the target storage system, and then migrate the data from the source storage system. As part of the data migration workflow, the host I/O paths are changed from the source storage system to the target storage system.
- After data migration
  The host can now use virtual information to access the resources on the target storage system. For example, when a Hitachi Universal Storage Platform V/VM is replaced by a VSP G1000, VSP G1500, or VSP F1500 and the information in Hitachi Universal Storage Platform V/VM is configured on a virtual storage machine, host configuration changes are not needed. The host can view and access the volumes on the VSP G1000, VSP G1500, or VSP F1500 virtual storage machine in the same way that the volumes were accessed on the Hitachi Universal Storage Platform V/VM.

For resources that have virtual information configured after data migration:

- User-defined virtual storage machine volumes cannot be allocated to a file server/NAS module.
- To manage copy pairs, such as global-active device pairs by using the resources of a virtual storage machine, review the prerequisites for system

configuration. For details, see the *Hitachi Command Suite Administrator Guide*.

Related concepts

- About resource groups on page 102
- About global-active device on page 334

Related tasks

- Creating virtual storage machines on page 327
- Editing virtual storage machines on page 330
- Deleting virtual storage machines on page 331

Related references

- Prerequisites for allocating resources to virtual storage machines on page 326
- Prerequisites for creating resource groups on page 104

Prerequisites for allocating resources to virtual storage machines

The following list identifies the conditions for resource allocation to virtual storage machines.

The following resources can be allocated to a virtual storage machine for VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, or VSP F1500.

- Parity Groups: Includes parity groups in external storage systems. When you register a part of a parity group that is part of a concatenated parity group to a virtual storage machine, other parity groups that are a part of the concatenated parity group will also be registered in the same virtual storage machine automatically.
- LDEV IDs: Includes volumes in external storage systems. LDEV IDs can be added only if they are not allocated to a host and do not have virtual information. Non-existent IDs can also be specified.
- Storage Ports: Includes storage system ports.
- Host Group/iSCSI Target Numbers: Host group/iSCSI target numbers can be added only if they are not allocated to a host and do not have virtual information. Non-existent numbers can also be specified.

Note the following additional conditions:

- If the LDEV ID of a DP volume is registered to a virtual storage machine, you can also view the DP pool to which the DP volume belongs and the DP pool volumes that compose the DP pool.

- If a DP pool volume is registered to a virtual storage machine, other volumes that compose the DP pool are automatically registered to the same virtual storage machine.
- If a volume allocation satisfies both of the following conditions, information about the volume might become incorrect, and users might be unable to complete tasks for volumes in Device Manager. Use the ID of a port supported by the model of the virtual storage machine for the ID of the physical port used for volume allocations. Alternatively, use Device Manager - Storage Navigator to complete tasks for allocated volumes:
  - The model of the virtual storage machine that is used for allocations differs from that of the physical storage system, and the ID of the physical port is not supported by the model of the virtual storage machine.
  - The Solaris multi-pathing feature (MPxIO) is enabled for the target host.



Note: In general, resources do not need to be added to a virtual storage machine before provisioning, GAD provisioning, or before performing nondisruptive migration. After resources are added to a virtual storage machine, they become unavailable to other virtual storage machines. In most cases, resources, such as volumes, will be added automatically to the virtual storage machine during volume allocation. Verify the requirements for using a virtual storage machine and add only those resources that are required.

Related concepts

- About virtual storage machines on page 322

Related tasks

- Creating virtual storage machines on page 327
- Editing virtual storage machines on page 330

Creating virtual storage machines

You can create a maximum of seven virtual storage machines for single VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500 or VSP F1500. Resources specified, such as parity groups or LDEV IDs, are added to the resource pools on the user-defined virtual storage machines.

Before you begin

Serial number, if you are specifying a storage system that is not registered with HCS.

Procedure

1. On the **Administration** tab, select **Virtual Storage Machine**.
2. In the application pane, click **Create Virtual Storage Machine**.

3. Enter a name for the virtual storage machine.
4. Choose one of the following options to specify the virtual model and serial number.
   - Click **Select a storage system already registered on HCS from list**.
   - Click **Specify manually**, select a model from the list, and then specify a virtual serial number.
5. On the **Storage Systems** tab, click **Add Storage Systems**, and then select the storage system to add.
   Displays the storage system and the resources in the storage system that can be added to the virtual storage machine. Resources are organized by resource type, such as parity groups, LDEV IDs, storage ports, and host group/iSCSI target numbers.
6. (Optional) On the **Parity Groups** tab, click **Add Parity Groups**, select the parity groups from the list, and then click **OK**.
7. (Optional) On the **LDEV IDs** tab, click **Add LDEV IDs**, specify the number of LDEV IDs or select available LDEV IDs to add from a table, and then click **OK**.
8. (Optional) On the **Storage Ports** tab, click **Add Storage Ports**, select the storage ports to add, and then click **OK**.
9. (Optional) On the **Host Group/iSCSI Target Numbers** tab, click **Add Host Group/iSCSI Target Numbers**, specify storage ports and the number of host group numbers/iSCSI targets to add per port, or specify host groups/iSCSI targets from a list, and then click **OK**.
10. Click **Submit.**
    If the task is scheduled to run immediately, the process begins.
11. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.

Result

The user-defined virtual storage machine is created and can be viewed in the list of virtual storage machines.

Related concepts

- About virtual storage machines on page 322

Related references

- Prerequisites for allocating resources to virtual storage machines on page 326

Allocating volumes to hosts by using virtual storage machine resources

The storage administrator can allocate volumes to hosts by using the resources of a virtual storage machine. The volumes of the default virtual

storage machine can also be allocated. The default virtual storage machine is created in VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, or VSP F1500.

Note: When you unallocate volumes, you can delete the virtual information and move the volume to the default virtual storage machine.

Before you begin

- Identify the name of the target host.
- The volumes to be allocated must satisfy one of the following conditions:
  - The volumes must belong to the same virtual storage machine to which the host group/iSCSI target of the target host belongs.
  - The volume must belong to the resource pool of the default virtual storage machine.

Procedure

1. On the **Resources** tab, select **Hosts**.
2. Expand the tree and select an operating system.
3. Select one or more hosts and click **Allocate Volumes**.
4. Select the **No. of Volumes** and enter the **Volume Capacity**.
5. From the list of storage systems, select the physical storage system.
6. From the list of virtual storage machines, select the virtual storage machine that contains the volumes to be allocated.

   Note: When allocating volumes from the default virtual storage machine, you do not need to select a virtual storage machine. The default virtual storage machine is selected by default.

7. Select the volume type and volume location.
8. (Optional) Click **Advanced Options** to specify tier, drive type, drive speed, RAID level, or parity group in the volume criteria. The storage system will search for volumes that match the criteria.
9. Click **LUN Path Options** to configure the number of LUN paths to allocate per host and assign or change LUN paths.
10. (Optional) Click **Host Group and LUN Settings** or **iSCSI Target and LUN Settings** to configure the host group/iSCSI target settings and LUN number.
11. Click **Show Plan** and confirm that the information in the plan summary is correct. If changes are required, click **Back**.
12. (Optional) Update the task name and provide a description.
13. (Optional) Expand **Schedule** to specify the task schedule.

    You can schedule the task to run immediately or later. The default setting is **Now**. If the task is scheduled to run immediately, you can select **View task status** to monitor the task after it is submitted.

14. Click **Submit.**
    If the task is scheduled to run immediately, the process begins.
15. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.
16. On the **Resources** tab, select hosts to verify that volumes have been allocated. The virtual information is set to the allocated volumes. When a default virtual storage machine volume is allocated, the volume is registered to the virtual storage machine to which the host groups/iSCSI targets belong.

Related references

- Prerequisites for allocating volumes on page 204

Editing virtual storage machines

You can add resources or remove resources from a virtual storage machine. You can also change the name of the virtual storage machine.

To support global-active device pairs, edit the virtual storage machine to add secondary storage system resources to a default or user-defined virtual storage machine on the primary storage system to provide a single view of the primary and secondary storage system resources.

Before you begin

From the host, unallocate volumes that are related to resources that belong to the virtual storage machine you are removing.

Procedure

1. On the **Administration** tab, select **Virtual Storage Machine**, select a virtual storage machine in the list, and then click **Edit Virtual Storage Machine**.
2. Verify the **Name** and **Virtual Model and Serial No.** of the virtual storage machine that you are modifying. If you are setting up a global-active device pair, verify that the virtual storage machine is on the primary storage system.
3. Add or remove resources, such as parity groups or LDEV IDs, for the virtual storage machine. If you are setting up a global-active device pair, click **Add Storage Systems**, select the secondary storage system, and then click **OK**. Then add secondary storage system resources, such as **Parity Groups**, **LDEV IDs**, **Storage Ports**, and **Host Group/iSCSI Target Numbers** to the virtual storage machine.
4. Enter an optional task description, select whether you want to view the progress of the task when the task completes, and then click **Submit**.
5. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.

Result

The selected virtual storage machine is edited and can be verified in the list of virtual storage machines. If you are creating a global-active device pair, you have a single view of the primary and secondary storage system resources.

Related concepts

- About virtual storage machines on page 322

Related references

- Prerequisites for allocating resources to virtual storage machines on page 326

Deleting virtual storage machines

You can delete virtual storage machines that are no longer required. Deleting a virtual storage machine also deletes all of the virtual information about the resources that belong to that virtual storage machine and removes the resources from the virtual storage machine.

Before you begin

From the host, unallocate the volumes that belong to the virtual storage machine that you are deleting.

Procedure

1. On the **Administration** tab, select **Virtual Storage Machine**.
2. From the list of virtual storage machines, select the virtual storage machine that you want to delete and click **Delete Virtual Storage Machine**.
3. Click **Submit**.
4. On the **Tasks & Alerts** tab, confirm the task completion.

Result

The selected virtual storage machine is deleted.

Related concepts

- About virtual storage machines on page 322

8

Replicating volumes using global-active device

Volume replication is used for creating in-system or remote storage system volume copies for data protection. Global-active device provides copy pair replication for active-active high availability in a multi-site environment.

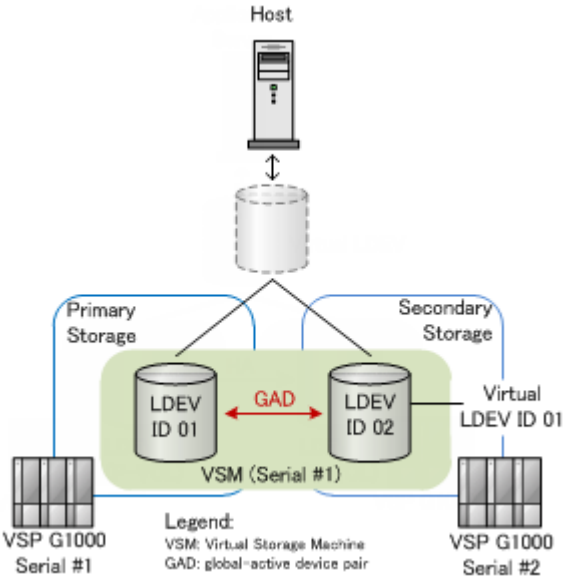
- ☐ About global-active device
- ☐ Setting up a global-active device environment
- ☐ Monitoring and managing global-active device pairs
- ☐ Recovering from global-active device failures
- ☐ Discontinuing a global-active device environment

About global-active device

Global-active device employs volume replication (P-VOL and S-VOL in a bidirectional copy pair) to provide a high availability environment for hosts across storage systems and sites. The purpose of global-active device is to provide data protection and minimize data access disruptions for host applications due to storage system or site failures.

A virtual storage machine represents storage resources such as parity groups and LDEVs (volumes). To support global-active device read/write operations, resources of the primary storage system (Serial #1) and secondary storage system (Serial #2) are defined in the virtual storage machine to provide a single view of the primary and secondary storage system resources.

As shown in the following figure, volumes on the secondary storage system are assigned a virtual LDEV ID that matches the physical LDEV ID of volumes belonging to the default virtual storage machine on the primary storage system. The host can then access the primary or secondary volume by using the same LDEV ID (01). This occurs when you create a global-active device pair during volume allocation.



VSP G1000, VSP G1500, VSP F1500, and VSP Gx00 models can be used for global-active device.

In addition to the example above, if the VSP G1000 microcode version in both the primary and secondary storage systems is `80-02-01-XX/XX` or later, and if the firmware version of VSP Gx00 models is `83-01-2X-XX/XX` or later, then you can also create a global-active device pair by setting, as the P-VOL, a volume that already has a virtual LDEV ID assigned and whose resources are managed by user-defined virtual storage machines. If you register a P-

VOL for a user-defined virtual storage machine, the virtual LDEV ID of volumes on the primary storage system are used.

Note: When you create a copy pair by using volumes with the capacity saving function enabled (dedupe and compression), copy or host I/O performance might be lowered. For details, see the manual for the applicable volume replication functionality.

The global-active device environment resolves the details of all read/write activity to the global-active device pair. For example, when the host writes data to LDEV ID 01, data is replicated to LDEV ID 02. If the host path to primary storage fails, the host can continue to perform read/write operations using the secondary storage volume.

When a global-active device pair is established, application server high-availability benefits include:

- Continuous I/O:
  If a primary volume becomes unavailable, the host will continue to transparently access the secondary volume.
- Clustered host failover:
  You do not need to perform storage system tasks, such as suspending or re-syncing a global-active device pair to recover from a host failure.
- Virtual machine migration:
  If a server is running a virtual machine, such as a VMware virtual machine running a guest OS, and creating a high load, you can move the virtual machine to the host at the other site, which can access the same data so that data migration between storage systems is unnecessary.

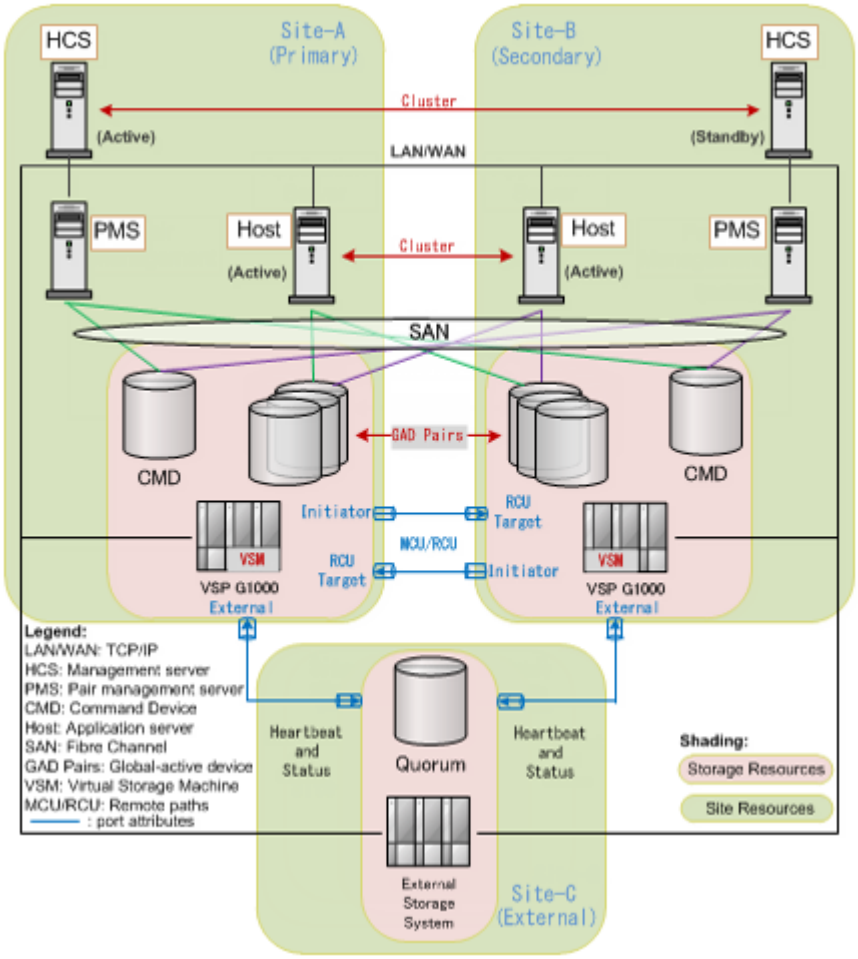
Using Replication Manager, configure in-system replication to further safeguard the data on global-active device pairs in use by hosts on primary and secondary storage systems.

The following configurations are supported:

- Cascade configuration
- Multi-target configuration
- GAD 3DC delta resync configuration

For more information about these configurations, see the *Hitachi Command Suite Replication Manager User Guide*.

A global-active device environment can be implemented within a single site or across two sites or three sites. Implementation across three sites provides the best protection against site failures because the primary, secondary, and quorum storage systems exist at different physical locations, or sites, as shown in the following figure.



In the figure above:

- The primary, secondary, and quorum storage systems are located across three sites, typically separated by short distances (campus) or intermediate (metro) distances. Distance can affect the required path management software for hosts. The primary and secondary storage systems are either VSP Gx00 models or VSP G1000, VSP G1500, and VSP F1500 storage systems. The quorum storage system monitoring the status of the primary and secondary storage can be any external storage system that supports virtualization of its volumes and is physically connected using Fibre Channel or Fibre Channel over Ethernet (FCoE) to 'External' ports configured on the primary and secondary storage systems. Supported external storage systems are identified, with their required port attributes, in the Hitachi Universal Volume Manager User Guide.
- The HCS management servers are clustered across two sites in an active-standby configuration. The active node is used to manage both sites. The standby node takes over site management if the active node fails. In addition to Device Manager and Tiered Storage Manager, Replication Manager is installed and licensed on both nodes for performing pair management tasks.

Note: You can set up and maintain GAD entirely using Device Manager:

- When using a single HCS management server node to discover the primary and secondary storage system, you can allocate global-active device pairs.
- Should a GAD pair failure occur, you can check the pair status and I/O mode using the Replication tab.
- You can also increase fault tolerance by changing a 2DC GAD configuration to 3DC delta resync, or adding local copy groups to existing GAD pairs.

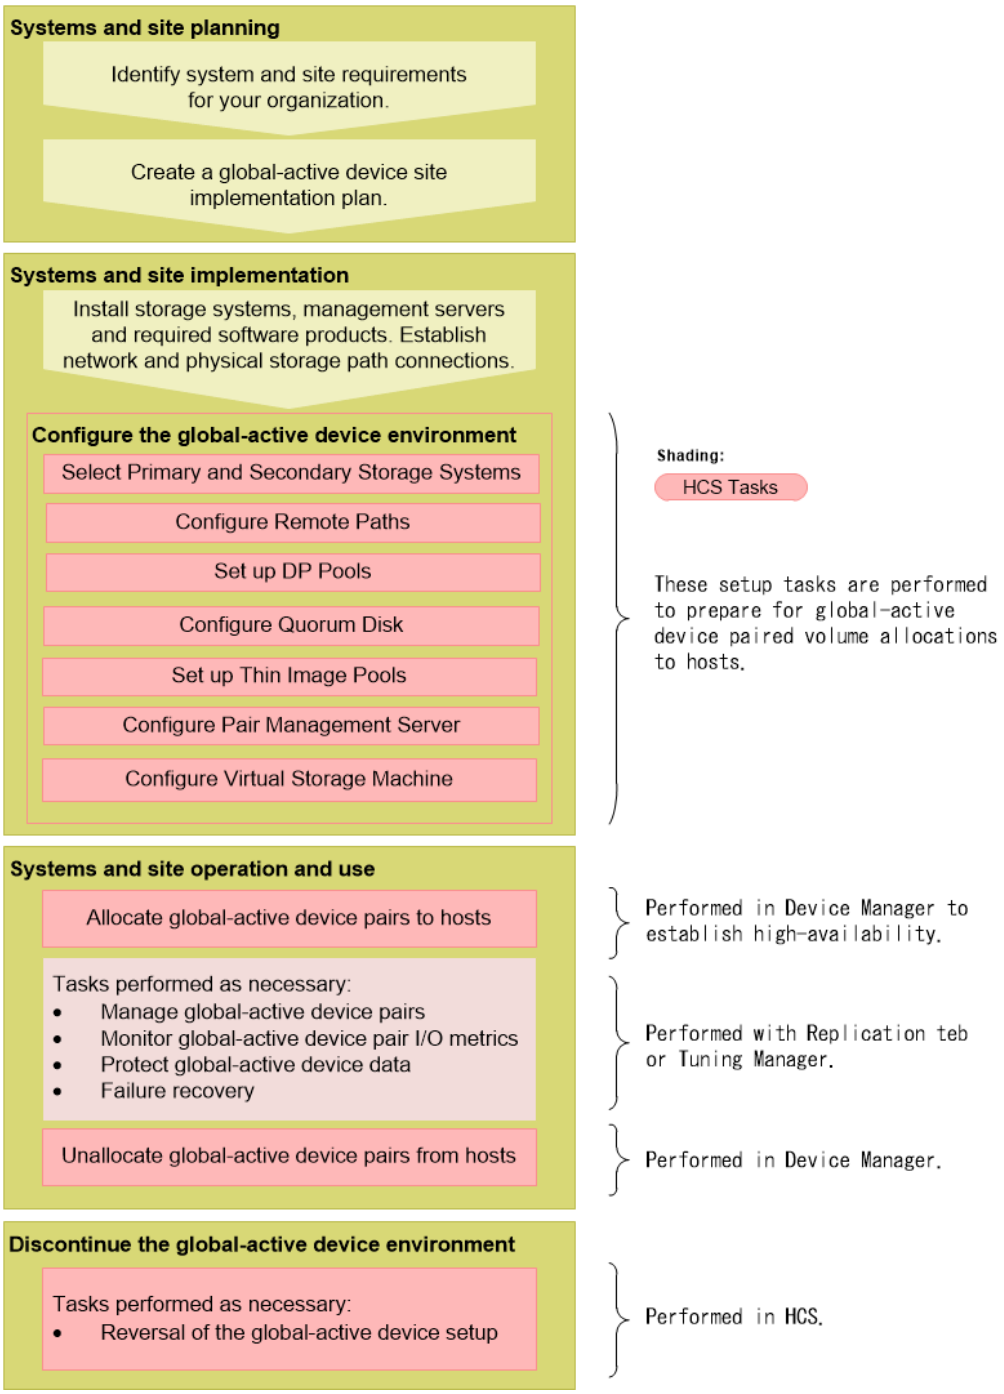
- Pair management servers are located at both the primary site and secondary site. Each pair management server must include Device Manager Agent and Command Control Interface (CCI) to perform pair management tasks as directed by HCS, or directly from the pair management server itself. Command devices (CMD) are used for processing pair management commands. The primary and secondary site pair management servers are used by the HCS management server for global-active device pair operations.
- The application servers are clustered across two sites with HBA (Fibre Channel) paths to both storage systems. Application servers can be configured in an active-standby or active-active cluster. In an active-active cluster configuration both nodes can perform read-write I/O. Depending on the multi-path software being used, an application server can use any available path to access data, or will use the shortest available path if configured to do so, meaning an application server at the primary site would read data from primary site storage, and an application server at the secondary site would read data from secondary site storage. Write I/O to global-active device pairs results in the same data being written to the pair, regardless of which storage system and port the write I/O occurs on. Data is always written to the primary volume first, then the secondary volume using the physical remote path (MCU/RCU) connections for bidirectional data transfers.

Note: Dynamic Link Manager (HDLM) provides features for identifying non-preferred (longer) paths for I/O, ensuring the use of the shortest I/O path under normal conditions. If a path failure occurs, the longer path is used for I/O. This optimizes I/O performance while retaining high-availability. See the *Command Suite Dynamic Link Manager User Guide* for your host OS. If you can suppress I/O on longer paths, then you can use ALUA, and the access efficiency improves. See the *Global-Active Device User Guide*. (Configuring ALUA is possible during the allocation of global-active device pairs in Device Manager.)

If the host OS is HP-UX 11iv3 and you are using the native multi-pathing features, execute the following command in advance to disable the multi-pathing features for legacy device special files:

```
scsimgr save_attr -a leg_mpath_enable=false
```

The following workflow describes the necessary planning and implementation steps for creating, monitoring and managing a global-active device environment.



For detailed information on site requirements, hardware setup, operational detail, and recovery recommendations, see the following:

- *Global-Active Device User Guide*
- *Provisioning Guide* for your storage system
- *Hitachi Command Suite Administrator Guide*
- *Hitachi Command Suite System Requirements*

Note: If an error occurs when you unallocate a global-active device pair volume or perform other operations on the pair volume, confirm the following:

- The Device Manager agent service is running on the pair management server.
- The command device allocated to the pair management server is working normally.

Finally, execute Refresh Hosts for the pair management server, and then perform the operation again.

Related concepts

- About global-active device setup on page 340

Setting up a global-active device environment

After the HCS management server has discovered resources, you can use the Set up Replication/GAD window for setup tasks that must be performed before global-active device pairs can be allocated.

About global-active device setup

The global-active device setup window provides a workflow for configuring required storage system resources before global-active device paired volumes can be allocated.

Access to the Set up Replication/GAD window is provided from:

- The Actions menu on the global task bar
- The Replication tab, from General Tasks

The setup window provides access to various HCS or Replication Manager dialog boxes. As tasks are completed, or if tasks have already been completed by other procedures, the setup window identifies completed tasks and provides a link to configuration details. If you close the setup window and reopen it later, the completed tasks remain identified.

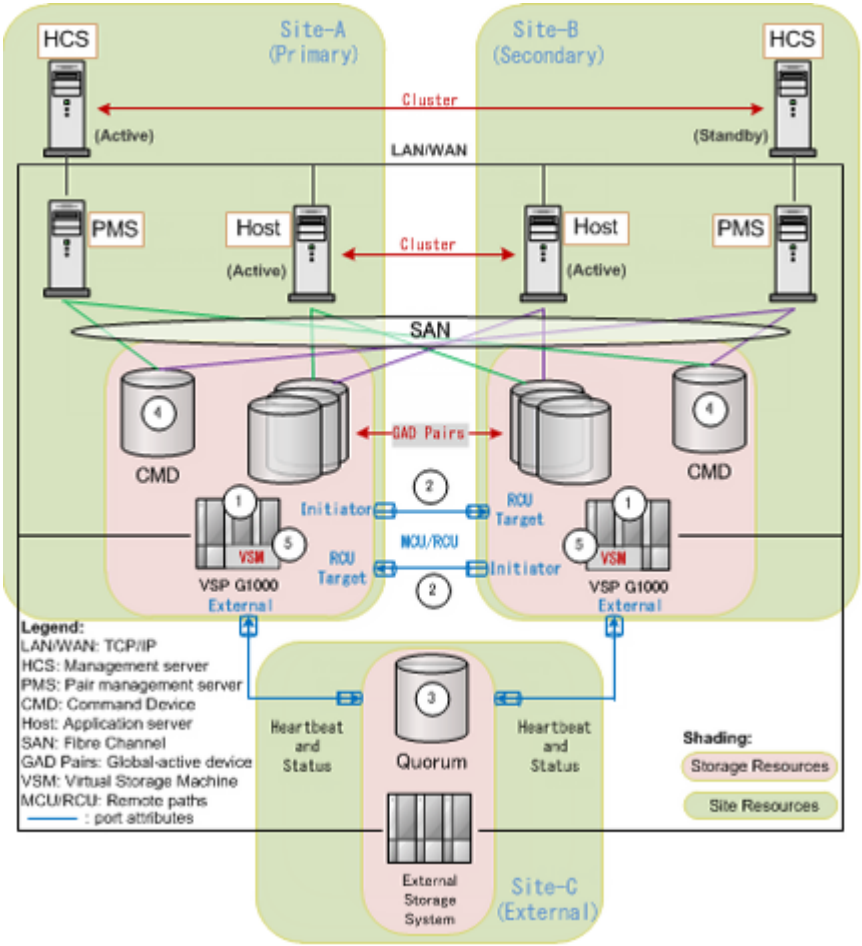
The following figure depicts the setup window configuration items in the global-active device environment:

- Select Storage Systems (see ① below, this task must be done first)

Note: When global-active device pairs are allocated, the primary storage system contains the P-VOLs and the secondary storage system contains the S-VOLs.

The following setup tasks can be performed in any order, but must be completed before allocating global-active device pairs.

- Configure Remote Paths (see ②)
- Set up DP pools (if you need to create DP pools for a command device, you must do so before configuring the pair management servers)
- Set up Thin Image pools
- Configure Quorum Disks (see ③)
- Configure Pair Management Servers (see ④)
- Configure Virtual Storage Machines (see ⑤)



Related concepts

- About global-active device on page 334

Related tasks

- Select primary and secondary storage systems on page 342
- Configuring remote paths on page 343
- Configuring the quorum disk on page 345
- Configuring pair management servers on page 348
- Configuring the virtual storage machine on page 349
- Creating DP pools for a replication configuration on page 295

- Creating Thin Image pools for a replication configuration on page 297

Related references

- Prerequisites for global-active device setup on page 342

Prerequisites for global-active device setup

Before performing global-active device setup tasks, ensure the following prerequisites are met:

- For VSP Gx00 models, VSP G1000, VSP G1500, and VSP F1500, the primary and secondary storage systems should be the same model. For VSP Gx00 models, the global-active device pair can be created even when the model numbers are different.
- The VSP G1000 microcode version for both the primary and secondary storage systems must be `80-02-01-XX/XX` or later, and the firmware version of VSP Gx00 models must be `83-01-2X-XX/XX` or later to create global-active device pairs where the P-VOLs already have a virtual LDEV ID assigned and are managed by user-defined virtual storage machines. This requirement is not necessary for global-active device pairs created by adding secondary storage resources to the primary storage default virtual storage machine.
- Global-active device has been licensed (which includes Replication Manager).
- Replication Manager is installed and linked to HCS.
- A pair management server is installed with Device Manager Agent and Command Control Interface (CCI) at the primary and secondary storage sites.
- The quorum storage system is connected to the primary and secondary storage system by Fibre Channel or Fibre Channel over Ethernet (FCoE).
- Port attribute settings for the primary and secondary storage system.
- You have registered (discovered) the primary, secondary, and quorum storage systems. If you are configuring quorum disks by using virtualized volumes as external volumes, registration of the quorum storage system is not necessary.

Related concepts

- About global-active device on page 334
- About global-active device setup on page 340

Related tasks

- Registering a storage system on page 60
- Creating virtual storage machines on page 327

Select primary and secondary storage systems

Select primary and secondary storage systems listed by the global-active device setup.

Procedure

1. On the **Replication** tab, **General Tasks**, select **Set up Replication/ GAD**.
2. Click **Select Storage Systems and Copy Type**.
3. From the **Select Storage Systems and Copy Type** window, select the primary and secondary storage systems from the list.
4. Select the copy type to be used between the storage systems.
5. Select a default or user-defined virtual storage machine on the primary storage system in which global-active device pairs will be created.
   (The drop-down list only displays the virtual storage machine of the primary storage system.)

   Note: To create global-active device pairs by using volumes in the virtual storage machine of the primary storage system for which a virtual LDEV ID has already been assigned, if the primary and secondary storage systems are both VSP G1000, the microcode version of each must be `80-02-01-XX/XX` or later. For VSP Gx00 models, the firmware version of each must be `83-01-2X-XX/XX` or later.

6. Confirm the selected storage systems and virtual storage machine information are correct and click **OK**.
   The selected information is displayed on the setup window as confirmation.

Result

The storage systems and virtual storage machine are selected, and become default values for the remaining setup tasks.

Related concepts

- About global-active device on page 334
- About global-active device setup on page 340

Configuring remote paths

Configure remote paths between the primary and secondary storage systems.

Before you begin

The primary and secondary storage systems have been selected.

Note: For more information about remote paths, see the *Hitachi Command Suite Replication Manager User Guide*.

Procedure

1. On the **Replication** tab, **General Tasks**, select **Set up Replication/ GAD**.
2. From **Set Up Global-Active Device**, click **Set Up Global-Active Device**.
3. From **Configure Remote Paths**, click **Create Remote Paths**.
   The **Create Remote Path** wizard opens.
4. On the **Select Remote Storage System** page:
   a. Verify that **Copy Type** is set to **global-active device**.
   b. Verify the displayed **Primary Storage System** and **Secondary Storage System** information, and click **Next** to define paths.

   Tip: The MCU/RCU check box is disabled by design for the global-active device dialog.

5. On the **Define Remote Path** page:
   a. Ensure that the **Select reverse direction path** check box is selected.
   b. (Optional) Enter a **Label** for identifying the remote paths for global-active device.
   c. (Optional) Enter a non-default **Path group ID**.
   d. Using the **Local Port** and **Remote Port** lists, select the Initiator and RCU target ports from the **Existing Ports** list for both the primary and secondary storage systems.

   Tip: The port list on the left shows primary to secondary storage ports, and the port list on the right shows secondary to primary storage ports. Note that the listed existing ports are unallocated host ports, meaning a port attribute of target not in use by any hosts.

   e. Click **Next** to review the remote path settings.
6. On the **Confirm** page:
   a. Under **Initiator/RCU Targets**, review selected ports and port attribute settings (before and after).
   b. In the **Remote Path** area, check the copy direction for the local and remote port pairs. If the port pairs are correct, select the confirmation message check box, and click **Confirm** to set the port attributes (Initiator or RCU Target).

   Tip: Port pairs on the left (Local -> Remote) are primary to secondary storage remote paths. Port pairs on the right (Local <- Remote) are the secondary to primary storage remote paths.

7. After performing the initial setup for a GAD environment, go to the **Set up Replication/GAD** window. From **Set Up Global-Active Device**, click **Refresh Display**. This changes the `Incomplete` status to `Complete`.

Result

The remote paths between the primary and secondary storage systems are configured, and can be confirmed by clicking Show Details in the Set up Replication/GAD window.

If setup of the remote path failed, in General Tasks, click Replication Manager. In the **Delete Remote Paths** dialog box, make sure the Do not change port attributes check box is selected, and then delete the remote path.

Related concepts

- About global-active device on page 334
- About global-active device setup on page 340

Configuring the quorum disk

The quorum disk detects the status of the primary and secondary storage systems. Configure the quorum disk by selecting a volume and virtualizing it as an external volume of the primary and secondary storage systems.

You can also configure the quorum disk by using a virtualized volume as an external volume of the primary or secondary storage system.

Tip: If you use the following storage systems when a quorum disk is blocked, the synchronization status of the copy pair data between the primary and secondary storage systems is detected, and the volume redundancy provided by global-active device is maintained:

- VSP G1000 (microcode version 80-05-0X/XX or later)
- VSP G1500
- VSP F1500
- VSP Gx00 models (microcode version 83-04-0X/XX or later)

Quorum disk configuration and supported storage systems

The following table shows whether the indicated storage system supports the quorum disk settings in the initial setup dialog box for the global-active device environment.

Task	VSP G1000		VSP G1500, VSP F1500	VSP Gx00 models
Quorum disk configuration	80-01-4X-XX/XX or later Earlier than 80-02-2X-XX/XX	80-02-2X-XX/XX or later	80-05-0X-XX/XX or later	83-01-2X-XX/XX or later

Task	VSP G1000		VSP G1500, VSP F1500	VSP Gx00 models
Create a new external volume*	Yes	Yes	Yes	Yes
Use an existing external volume	No	Yes	Yes	Yes

*The external path using iSCSI is not supported.

Before you begin

- The primary and secondary storage systems have been selected.
- You can create a quorum disk by using an external volume that connects using iSCSI. However, the HCS global-active device environment supports only displaying quorum disk information and verifying the configuration, not creating a quorum disk.

  To specify an iSCSI-connected external volume for the quorum disk, first create an external volume by clicking the System GUI link and creating an external volume that connects using iSCSI. Then, you can select the external volume you created in Configure Quorum Disks in the initial setup dialog box of the HCS global-active device workflow.

Procedure

1. On the **Replication** tab, **General Tasks**, select **Set up Replication/ GAD**.
2. From **Set Up Global-Active Device**, click **Set Up Global-Active Device**.
3. From **Configure Quorum Disks**, when the firmware version of VSP Gx00 models for both the primary and secondary storage systems is 83-01-2X-XX/XX or later and when the VSP G1000 microcode version for both the primary and secondary storage systems is 80-02-2X-XX/XX or later, do one of the following:
   - Click **Create a new external volume** to virtualize a volume as an external volume, and then continue with step 3.
   - Click **Configure an existing external volume** to use a virtualized volume as an external volume, and then click **Configure Quorum Disks** for the **Primary Site** and **Secondary Site**. On the **Quorum Disk** tab, click **Add Quorum Disks**, and then follow the instructions in the Help.
4. In the **Select Volume** window, select a **Storage system** to provide the quorum disk, and then do one of the following:
   - From the **Available Volumes** list, select a volume, click **OK**, and then go to step 12.
   - If no suitable volume is listed, click **Create Volumes**.

5. In the **Create Volumes** window, create a volume to be used as the quorum disk.
6. From the **Volume Type** list select a volume type.
7. (Optional) Click **Advanced Options** to configure volume options or specify a quorum disk label.
8. Click **Show Plan** and confirm that the information in the plan summary is correct. If changes are required, click **Back**.
9. (Optional) Update the task name and provide a description.
10. (Optional) Expand **Schedule** to specify the task schedule.

    You can schedule the task to run immediately or later. The default setting is **Now**.
11. Click **Submit.**
    If the task is scheduled to run immediately, the process begins.
12. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.
13. From **Configure Quorum Disks**, for the primary storage, click **Virtualize Volumes**.
14. In the **Virtualize Volumes** dialog box:
    a. Verify the external and internal storage systems are correct.
    b. Click **Edit External Paths** to set or change external paths.
    c. (Optional) Select **External Path Priority Settings** to increase or decrease path priority.
    d. Select **Internal Volume Configuration** and verify that **Change Internal Volume(s) to Quorum Disk(s)** is selected by default. The default ID is recommended, but can be changed to another value if necessary.
15. Click **Show Plan** and submit the task.
16. Repeat steps 12 through 14 for the secondary storage system.

    Note that for the secondary storage system, the secondary storage system name will be displayed.
17. After performing the initial setup for a GAD environment, go to the **Set up Replication/GAD** window. From **Set Up Global-Active Device**, click **Refresh Display**. This changes the `Incomplete` status to `Complete`.

Result

The quorum disk is created and virtualized for both the primary and secondary storage, and can be confirmed by clicking Show Details in the Set up Replication/GAD window.

Related concepts

- About global-active device on page 334
- About global-active device setup on page 340

Related references

- Create Volumes dialog box on page 152
- Virtualize Volumes dialog box on page 166

Configuring pair management servers

Configure the command devices for the pair management servers from the primary and secondary storage systems. Command devices are used to communicate pair management directives.

Before you begin

- The primary and secondary storage systems have been selected.
- Install CCI and Device Manager agent on the pair management server. (When you install a Device Manager agent on a pair management server, the pair management server will be registered as a host in HCS.)

Note: If you are using a virtual machine as a pair management server and the virtualization server uses NPIV (N-Port ID Virtualization) HBAs, you must also register the virtualization server in HCS.

Procedure

1. On the **Replication** tab, **General Tasks**, select **Set up Replication/ GAD**.
2. From **Set Up Global-Active Device**, click **Set Up Global-Active Device**.
3. From **Configure Pair Management Server**, click **Allocate Volumes** for the primary storage system.
4. In the **Allocate Volumes** dialog box, verify that the **Storage System** is correct.
5. For **Host**, select the pair management server for the storage system.
6. Click **Advanced Options**, verify that **Command Device** is checked, and that **User Authentication** is **Enabled**.
7. Click **Show Plan** and confirm that the information in the plan summary is correct. If changes are required, click **Back**.
8. (Optional) Update the task name and provide a description.
9. (Optional) Expand **Schedule** to specify the task schedule.

   You can schedule the task to run immediately or later. The default setting is **Now**.
10. Click **Submit.**
    If the task is scheduled to run immediately, the process begins.
11. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.
12. Repeat steps 2 through 10 for the secondary storage system. For **Host**, be sure to specify the pair management server for the secondary storage system.

13. Ensure that the pair management servers recognize the command devices, and then from **Configure Pair Management Server**, click **Refresh Hosts**.
14. After performing the initial setup for a GAD environment, go to the **Set up Replication/GAD** window. From **Set Up Global-Active Device**, click **Refresh Display**. This changes the `Incomplete` status to `Complete`.

Result

The command devices for the primary and secondary pair management servers are allocated, and pair management server information is current, and can be confirmed by clicking Show Details in the Set up Replication/GAD window.

Related concepts

- About global-active device on page 334
- About global-active device setup on page 340

Related references

- Allocate Volumes dialog box on page 213

Configuring the virtual storage machine

To support global-active device pairs, add secondary storage system resources to the default or user-defined virtual storage machine on the primary storage system. Configuring the virtual storage machine provides a single view of the primary and secondary storage system resources.

Before you begin

- The primary and secondary storage systems have been selected.
- Identify the primary and secondary storage system resources to be added to the virtual storage machine.

Note: For the host group or iSCSI target number, use a host group or iSCSI target number that has not been allocated to a storage port.

Procedure

1. On the **Replication** tab, **General Tasks**, select **Set up Replication/GAD**.
2. From **Set Up Global-Active Device**, click **Set Up Global-Active Device**.
3. From **Configure Virtual Storage Machine**, click **Edit Virtual Storage Machine**.
4. In the **Edit Virtual Storage Machine** dialog box, on the **Storage Systems** tab, verify that the displayed **Name** and **Virtual Model and**

Serial No. show the primary storage system values and that the primary storage system is listed.

5. Click **Add Storage Systems**, select the secondary storage system, and then click **OK**.
6. Add secondary storage system resources, such as **Parity Groups**, **LDEV IDs**, **Storage Ports** and **Host Group/iSCSI Target Numbers** to the virtual storage machine by clicking the respective tabs.

   Note: The primary storage system resources already belong to the virtual storage machine.

7. When you are finished adding resources from the secondary storage system, enter an optional task description, select whether you want to view the progress of the task, and then click **Submit**.
8. After performing the initial setup for a GAD environment, go to the **Set up Replication/GAD** window. From **Set Up Global-Active Device**, click **Refresh Display**. This changes the `Incomplete` status to `Complete`.

Result

The storage system resources are configured for the virtual storage machine.

Related concepts

- About global-active device on page 334
- About global-active device setup on page 340

Related tasks

- Creating virtual storage machines on page 327
- Editing virtual storage machines on page 330

Monitoring and managing global-active device pairs

This module describes the required tasks to allocate, monitor, and unallocate global-active device pairs for hosts.

Allocating global-active device pairs

Allocate new global-active device pairs to hosts (application servers). When you allocate global-active device pairs, an initial copy is performed for the primary and secondary volumes.

Before you begin

- You must be authenticated to the storage system and logged in to the pair management server. For details, see the *Hitachi Command Suite Administrator Guide*.
- When you assign a CTG ID for copy pairs, make sure that the same CTG ID is assigned for the copy pairs within the copy group. You cannot use

CTG IDs used by another copy group. CTG IDs need to be assigned if global-active device copy pairs are used in a 3DC delta resync configuration that includes Universal Replicator.

Procedure

1. On the **Resources** tab, **General Tasks**, select **Allocate Volumes**.

   Tip: If you do not see **Allocate Volumes** listed, click **more...** to see all menu items.

2. In the **Allocate Volumes** dialog box, select the hosts for the global-active device pair.

   Note: File servers/NAS modules are not supported for global-active device pair allocations.

3. For **Allocation Type** select **global-active device**.
4. Specify the number and capacity of the volumes.
5. On the **Primary** tab, configure the following options:
   a. Select the primary **Storage system**, **Virtual Storage Machine**, and **Volume type**.
   b. Click **Advanced Options** for further configuration.
   c. In **LUN Path Options**:
      Edit LUN paths from primary storage ports to hosts.
   d. Click **Host Group and LUN settings** for further configuration.

      Note: If the host operating system is Windows or AIX and you are using the MPIO function, set the host mode option of the host group to 102.

   e. In **Pair management Server Settings**:
      Select an available host as the pair management server. For Instance ID, click **Existing** to use a current configuration definition file and related instance ID number on the pair management server, or click **New** to create a new configuration definition file, and enter a new instance ID number which will be validated as an available instance ID. You can also specify a **UDP Port** number for the instance.
6. On the **Secondary** tab, configure the following options:
   a. Select the secondary **Storage system**, **Virtual Storage Machine**, and **Volume type**.
      Virtual storage machine information defaults to the primary tab value.
   b. In **Advanced options**:
      When you select Basic volumes on the Primary tab, the volumes that correspond to the settings specified on the Primary tab, are displayed

on the Secondary tab. Unlike the Primary tab, the Secondary tab allows you to select secondary storage system DP pools (Pools for S-VOL) that correspond to primary storage system DP pools configured on the primary tab (Pools for P-VOL). For example, automatic selection of DP volumes will display primary DP pools, and eligible secondary DP pools for volume allocation. You can examine primary tab DP pool and volume details with links. You can accept the DP pools selected for secondary volumes, or make a manual selection by clicking **Select Pool** to list and select eligible DP pools, and selecting eligible volumes from the DP pools.

c. In **LUN Path Options**:

   Edit LUN paths from secondary storage ports to hosts. Under **LUN Path Setting Information on Primary Site**, you will see the primary storage system, and No. of LUN Paths per volume. Clicking the link will display primary storage system mapping information.

d. Click **Host Group and LUN settings** for further configuration.



Note: If the host operating system is Windows or AIX and you are using the MPIO function, set the host mode option of the host group to 102.

e. In **Pair management Server Settings**:

   Select an available host as the pair management server. For Instance ID, click **Existing** to use a current configuration definition file and related instance ID number on the pair management server, or click **New** to create a new configuration definition file, and enter a new instance ID number which will be validated as an available instance ID. You can also specify a **UDP Port** number for the instance.

7. Click **Virtual ID Settings** for further configuration.

These settings are displayed only in a configuration where options other than the default virtual storage machine are used on the primary side.

8. In **Pair Settings**:

a. Select the **Quorum Disk**, and then select whether pairs should belong to a new or existing **Copy Group** for management purposes. Click **New** to assign a new copy group to the global-active device pair. Select the **CTG ID** check box to assign a consistency group (CTG) ID to the new copy group. Click **Existing** to assign an existing copy group to the global-active device pair. The existing CTG ID will be used if it exists for the selected copy group.



Tip: To add the GAD 3DC delta resync configuration to allocated global-active device pair volumes, CTG IDs need to be assigned.

b. Manually enter a name (Prefix and Start Sequence No.) or select **Automatic** for **Pair Name** and a name will be created.

9. Click **Show Plan** and confirm that the information in the plan summary is correct. If changes are required, click **Back**.

10. (Optional) Update the task name and provide a description.
11. (Optional) Expand **Schedule** to specify the task schedule.
    You can schedule the task to run immediately or later. The default setting is **Now**.
12. Click **Submit.**
    If the task is scheduled to run immediately, the process begins.
13. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.
14. Confirm the global-active device pair is allocated.



Note: If you change the ALUA attribute of a volume for which a LUN path is set, you must rediscover the volume on the host.

Result

The global-active device pair is allocated to the target host, and initial copy performed. For global-active device pairs that use DP pool volumes, reclaiming zero pages is also automatically performed.

Related tasks

- Allocating global-active device pairs based on like volumes on page 353
- Allocating global-active device pairs using allocated volumes on page 356
- Editing LUN paths on page 240

Related references

- Allocate Volumes dialog box on page 213

Allocating global-active device pairs based on like volumes

Select an existing allocated global-active device volume, and allocate a global-active device pair using the characteristics of the selected primary or secondary volume.

Before you begin

You must be authenticated to the storage system and logged in to the pair management server. For details, see the *Hitachi Command Suite Administrator Guide*.

Procedure

1. On the **Resources** tab, select a global-active device volume from the **Open-Allocated** volume list, or for a specific host using global-active device paired volumes, select an already allocated volume to use in a new global-active device pair, and click **Allocate Like Volumes**.

Tip: Regardless of whether the primary volume (P-VOL) or the secondary volume (S-VOL) is selected as the existing volume, the settings for the existing P-VOL are used for the newly allocated primary volume and the settings for the existing S-VOL are used for the newly allocated secondary volume.

However, if characteristics such as drive type and drive speed of the existing volumes do not match between the primary and secondary storage systems, a secondary volume whose characteristics match those of the primary volume is recommended.

2. In the **Allocate Like Volumes** dialog box, verify the information for the selected volume and make any necessary changes.
3. On the **Primary** tab, configure the following options:
   a. Select the primary **Storage system**, **Virtual Storage Machine**, and **Volume type**.
   b. Click **Advanced Options** for further configuration.
   c. In **LUN Path Options**:

      Edit LUN paths from primary storage ports to hosts.
   d. Click **Host Group and LUN settings** for further configuration.

      Note: If the host operating system is Windows or AIX and you are using the MPIO function, set the host mode option of the host group to 102.

   e. In **Pair Management Server Settings**:

      Select an available host as the pair management server. For Instance ID, click **Existing** to use a current configuration definition file and related instance ID number on the pair management server, or click **New** to create a new configuration definition file, and enter a new instance ID number which will be validated as an available instance ID. You can also specify a **UDP Port** number for the instance.
4. On the **Secondary** tab, configure the following options:
   a. Select the secondary **Storage system**, **Virtual Storage Machine**, and **Volume type**.

      Virtual storage machine information defaults to the primary tab value.
   b. In **Advanced Options**:

      When you select Basic volumes on the Primary tab, the volumes that correspond to the settings specified on the Primary tab, are displayed on the Secondary tab. Unlike the Primary tab, select DP pools for the secondary storage system (Pools for S-VOL) that correspond to the DP pools configured on the primary tab (Pools for P-VOL). For example, automatic selection of DP volumes will display primary DP pools, and eligible secondary DP pools for volume allocation. You can examine primary tab DP pool and volume details with links. You can accept the

DP pools selected for secondary volumes, or make a manual selection by clicking **Select Pool** to list and select eligible DP pools, and selecting eligible volumes from the DP pools.

c. In **LUN Path options**:

Edit LUN paths from secondary storage ports to hosts. **LUN Path Setting Information on Primary Site** lists the primary storage system, and No. of LUN Paths per volume. Click the link to display primary storage system mapping information.

d. Click **Host Group and LUN settings** for further configuration.

Note: If the host operating system is Windows or AIX and you are using the MPIO function, set the host mode option of the host group to 102.

e. Click **Pair Management Server Settings** for further configuration.

5. Click **Virtual ID Settings** for further configuration.

   These settings display only when options other than the default virtual storage machine are used on the primary storage.

6. In **Pair Settings**:

   a. Select the **Quorum Disk**, and then select whether pairs should belong to a new or existing **Copy Group** for management purposes. Click **New** to assign a new copy group to the global-active device pair. Select the **CTG ID** check box to assign a consistency group (CTG) ID to the new copy group. Click **Existing** to assign an existing copy group to the global-active device pair. The existing CTG ID will be used if it exists for the selected copy group.

   Tip: To add the GAD 3DC delta resync configuration to allocated global-active device pair volumes, CTG IDs need to be assigned.

   b. Manually enter a name (Prefix and Start Sequence No.) or select Automatic for **Pair Name** and a name will be created.

7. Click **Show Plan** and confirm that the information in the plan summary is correct. If changes are required, click **Back**.
8. (Optional) Update the task name and provide a description.
9. (Optional) Expand **Schedule** to specify the task schedule.

   You can schedule the task to run immediately or later. The default setting is **Now**.

10. Click **Submit.**
    If the task is scheduled to run immediately, the process begins.
11. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.
12. Confirm the global-active device paired volumes are allocated.

Result

The global-active device pair is allocated to the target host, and initial copy performed. For global-active device pairs that use DP pool volumes, reclaiming zero pages is also automatically performed.

Related tasks

- Allocating global-active device pairs on page 350
- Allocating global-active device pairs using allocated volumes on page 356
- Editing LUN paths on page 240

Related references

- Allocate Volumes dialog box on page 213

Allocating global-active device pairs using allocated volumes

Existing open-allocated volumes in the same storage system belonging to the same host group (with the same LUN paths) can be converted to P-VOLs, and the P-VOL data migrated to S-VOLs to create a global-active device pair.

Before you begin

You must be authenticated to the storage system and logged in to the pair management server. For details, see the *Hitachi Command Suite Administrator Guide*.

Procedure

1. On the **Resources** tab, or by searching, locate allocated volumes to convert into global-active device pairs.
2. Select the volume, and click **Change to Global-Active Device Volumes**.
3. On the **Primary** tab, configure the following options:
   a. Verify the displayed hosts, storage systems virtual storage machines for the selected volumes.
   b. **Selected Volumes** are allocated volumes eligible to become a P-VOL in a global-active device pair. Details such as WWN/iSCSI Name, host group/iSCSI target, and volume characteristics are displayed.
   c. In **Pair Management Server Settings**:

      Select an available host as the pair management server. For Instance ID, select **Existing** to use a current configuration definition file and related instance ID number on the pair management server, or select **New** to create a new configuration definition file, and enter a new instance ID number which will be validated as an available instance ID. You can also specify a **UDP Port** number for the instance.
4. On the **Secondary** tab, configure the following options:
   a. Verify the displayed hosts for the selected volumes, and click **Select other Hosts** to specify more hosts that can access global-active device pairs.

b. Select the secondary storage system where the S-VOLs will be created.
c. Verify the displayed virtual storage machine.
d. Click **Advanced Options** for further configuration.
e. Click **LUN Path Options** for further configuration.
The previously allocated primary path is displayed.
f. Click **Host Group and LUN settings** for further configuration.



Note: If the host operating system is Windows or AIX and you are using the MPIO function, set the host mode option of the host group to 102.

g. In **Pair Management Server Settings**:
Select an available host as the pair management server. For Instance ID, select **Existing** to use a current configuration definition file and related instance ID number on the pair management server, or select **New** to create a new configuration definition file, and enter a new instance ID number which will be validated as an available instance ID. You can also specify a **UDP Port** number for the instance.

5. In **Pair Settings**:
a. Select the **Quorum Disk**, and then select whether pairs should belong to a new or existing **Copy Group** for management purposes. Click **New** to assign a new copy group to the global-active device pair. Select the **CTG ID** check box to assign a consistency group (CTG) ID to the new copy group. Click **Existing** to assign an existing copy group to the global-active device pair. The existing CTG ID will be used if it exists for the selected copy group.



Tip: To add the GAD 3DC delta resync configuration to allocated global-active device pair volumes, CTG IDs need to be assigned.

b. Manually enter a name (Prefix and Start Sequence No.) or select Automatic for **Pair Name** and a name will be created.
6. Click **Show Plan** and confirm that the information in the plan summary is correct. If changes are required, click **Back**.
7. (Optional) Update the task name and provide a description.
8. (Optional) Expand **Schedule** to specify the task schedule.
You can schedule the task to run immediately or later. The default setting is **Now**.
9. Click **Submit.**
If the task is scheduled to run immediately, the process begins.
10. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.
11. Confirm the global-active device paired volumes are created by checking volume status information.

Note: If you change the ALUA attribute of a volume for which a LUN path is set, you must rediscover the volume on the host.

Result

The global-active device pair has been created from open-allocated volumes, and initial copy performed. For global-active device pairs that use DP pool volumes, reclaiming zero pages is also automatically performed.

Related tasks

- Allocating global-active device pairs on page 350
- Allocating global-active device pairs based on like volumes on page 353
- Editing LUN paths on page 240

Related references

- Allocate Volumes dialog box on page 213

Checking global-active device pair status

You can review information (including status) of global-active device pairs or reserved volumes from the volume list and volume details windows in HCS.

Note: For information on the meaning of the pair status display, see the *Replication Manager Online Help* or the *Hitachi Command Suite Replication Manager User Guide*.

Procedure

1. Open the **Replication** tab.
2. If you have received an alert, go the **Alerts** table and find the entry that matches the error received.

   If you are simply checking the status of a replication configuration, go to the **Copy Topology** table.
3. Click the link for the copy topology of the global-active device you wish to examine.
4. From the **Copy Groups** or **Copy Pairs** tables, locate and select the entry showing an error.
5. (Optional) **Click Change Pair Status** to correct the status of the copy group or copy pair.
6. (Optional) In the application area, click the **View HRpM Tasks** link to confirm the task has completed.

Related tasks

- Monitoring global-active device pair performance on page 359

Monitoring global-active device pair performance

Hitachi Tuning Manager will collect and display performance information for global-active device pairs. If an administrator suspects a performance problem, or receives an alert indicating a performance problem, the administrator can investigate and resolve the issue.

Because global-active device pairs are kept synchronized, I/O issues on either the primary or secondary storage system can impact host performance. Status information and I/O metrics related to P-VOL and S-VOL performance can be examined to identify the problem, and resolve host I/O issues.

Before you begin

- Hitachi Tuning Manager must be installed, licensed and configured.
- When selecting one member of a global-active device pair, a single Device Manager managing the pair is required for analyzing the performance of the P-VOL and S-VOL simultaneously.

Procedure

1. On the **Analytics** tab, open the **Hosts** tree, select a host OS, and then select the target host as follows:
   a. Select the radio button next to the target host name, and then click **Identify Performance Problems**.
   b. Click **Host Name** to select target host volumes, and then click **Identify Performance Problems**.

   Tip: When you select volumes from a host volume list or from within the **Identify Performance Problems** wizard, if you select a global-active device volume in a single Device Manager environment, the P-VOL and S-VOL will both be targets for analysis. To identify global-active device volumes, use **Column Settings** to include **Copy Info (P-VOL)**, **Copy Info (S-VOL)**, and **Virtual LDEV ID** in the host volume list or the Identify Performance Problems wizard.

   The **Identify Performance Problems** wizard is displayed (starting on the Introduction page). Review displayed messages.
2. Click **Next**.
3. On the **Target Selection** page:
   a. If prompted, select the **Storage system type/volume type (location)**. For example: {VSP G1000, VSP G1500, VSP F1500} / DP Volume (Internal).
   b. Click **List of added volumes** to view analysis targets, including global-active device volumes.

Note: When the global-active device pair (P-VOL and S-VOL) of a selected global-active device volume are not found in HCS, the volume cannot be set to a target for analysis. These volumes are listed separately.

c. In **Select volumes to analyze**, **Select all** selects all listed volumes for analysis. Click **Select manually** to individually select volumes or global-active device pairs (P-VOL and S-VOL) for analysis, and then click **Add** to update the **Selected Volumes** list.
d. For **Select time period to analyze within 30 days before the current time**, configure the time period to analyze.
e. Click **Next**.

4. On the **Overview** page:
   a. Check the analysis results and cautions. If there is a resource that exceeds a threshold value, there might be a performance issue related to the storage system.
   b. Click **Next**.
5. On the **Volume** page:
   a. Review the **Response Time** and **IOPS** graphs displaying both P-VOL and S-VOL performance metrics.

Tip: If there is a large disparity between the performance metrics for the P-VOL and S-VOL, complete the wizard steps to investigate the issue further and determine the root cause. For example, there could be significant load variances between the connected hosts for the primary/secondary storage systems, remote path latencies, or multi-path software configuration issues.

6. Continue stepping through the **Identify Performance Problems** wizard to review further information on Cache, MP Blades, and Parity Groups.

Result

The performance metrics of global-active device pairs have been displayed for analysis.

Related concepts

- About analyzing storage system performance on page 526

Related tasks

- Checking global-active device pair status on page 358

Unallocating global-active device pairs

When a host with global-active device pairs is no longer needed, or the volume redundancy provided by global-active device pairs is no longer required, you can unallocate the global-active device pairs.

Depending on conditions, when a user unallocates global-active device volumes, the copy pair relationship will be ended, and related copy groups and configuration definition files (`horcmN.conf`) could be deleted as a result.

Note: When you unallocate the primary volumes of non-default virtual storage machines, an error might result. If this occurs, check the volume status and then delete virtual information and the GAD Reserved attribute. To delete this information, use Device Manager - Storage Navigator or the **`ModifyVLDEV`** and **`CancelVLDEVReservation`** command of the Device Manager CLI. After doing so, unallocate the volumes.

Before you begin

- Identify the name of the target host, and the global-active device volumes to unallocate.
- If necessary, back up data on the target volumes.
- Unmount all allocated volumes that you plan to unallocate. An IT administrator might have to perform this task.

Procedure

1. Select a host OS, then select a host using global-active device volumes and click **Unallocate Volumes**.

   Tip: If you select specific volumes for the host, **Unallocate Volumes** will go directly to the **Show Plan** window.

2. From the **Unallocate Volumes** dialog box, verify that the host is correct, and then select the P-VOL or S-VOL of the global-active device pair, and select **Unallocate global-active device pair simultaneously**.
3. Click **Show Plan** and confirm that the information in the plan summary is correct. If changes are required, click **Back**.
4. (Optional) Update the task name and provide a description.
5. (Optional) Expand **Schedule** to specify the task schedule.

   You can schedule the task to run immediately or later. The default setting is **Now**. If the task is scheduled to run immediately, you can select **View task status** to monitor the task after it is submitted.
6. Click **Submit.**

   If the task is scheduled to run immediately, the process begins.
7. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.

Result

Unallocated global-active device volumes are added back to the storage system Open-Unallocated volume list.

Related concepts

- About removing hosts and related resources, or vCenters on page 79

Related tasks

- Unallocating individual global-active device volumes on page 362
- Discontinuing a global-active device environment on page 365

Related references

- Unallocate volumes dialog box on page 229

Unallocating individual global-active device volumes

Individual global-active device primary or secondary volumes can be unallocated.

Before you begin

- Identify the name of the target host, and the global-active device primary or secondary volume to unallocate.
- If necessary, back-up data for the target volume.

Note:

- If you want to delete copy pairs for which CTG IDs were assigned and continue I/O with the secondary volumes, you need to perform operations on a per-copy-group basis. Delete all copy pairs within the copy topology from the Replication tab, or use the `DeleteReplication` command of the Device Manager CLI to delete all copy pairs within the copy group to which the target copy pair belongs.
- When you unallocate the primary volumes of non-default virtual storage machines, an error might result. If this occurs, check the volume status and then delete virtual information and the GAD Reserved attribute. To delete this information, use Device Manager - Storage Navigator or the `ModifyVLDEV` and `CancelVLDEVReservation` command of the Device Manager CLI.

After performing either of these operations, be sure and unallocate the volumes.

Procedure

1. On the **Resources** tab, select a host OS, then select a host using global-active device volumes, and click **Unallocate Volumes**.

   Tip: If you select specific volumes for the host, **Unallocate Volumes** will go directly to the **Show Plan** window.

2. In the **Unallocate Volumes** dialog box, verify that the host is correct, and then select the P-VOL or S-VOL of the global-active device pair to unallocate.

   Note: When unallocating the secondary volume, the primary volume is retained as a basic volume. When unallocating the primary volume, the secondary volume is retained as a basic volume.

3. Click **Show Plan** and confirm that the information in the plan summary is correct. If changes are required, click **Back**.
4. (Optional) Update the task name and provide a description.
5. (Optional) Expand **Schedule** to specify the task schedule.

   You can schedule the task to run immediately or later. The default setting is **Now**. If the task is scheduled to run immediately, you can select **View task status** to monitor the task after it is submitted.
6. Click **Submit.**
   If the task is scheduled to run immediately, the process begins.
7. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.

Result

Unallocated global-active device primary or secondary volumes are added back to the storage system Open-Unallocated volume list.

Related concepts

- About removing hosts and related resources, or vCenters on page 79

Related tasks

- Unallocating global-active device pairs on page 360
- Discontinuing a global-active device environment on page 365

Related references

- Unallocate volumes dialog box on page 229

Recovering from global-active device failures

This module provides an overview of global-active device failures, where to find detailed information, and specific tasks that can be performed from HCS.

About global-active device failures

Failures need to be identified when they occur, and actions taken to remedy the failure and re-synchronize global-active device pairs.

Failure types include:

- Host path failures (HBAs, SAN switches, or Target ports) between hosts and primary or secondary storage systems.
- Remote path failures (Initiator/RCU Target ports) between primary and secondary storage systems.
- External storage (quorum) path failures between the external storage system and primary/secondary storage systems, or failure of the quorum disk, or failure of the quorum storage system.
- Primary or Secondary storage system failure.
- Primary or Secondary site failure.

Failures typically cause global-active device pair status to change from normal (for example, PAIR status) to some combination of status information indicating unavailability of a P-VOL, an S-VOL, or both. After recovery of the global-active device pair, you can verify if the I/O mode of an S-VOL is changed to mirror status from the Replication tab.

For all failures (except host path failures) SIM reference codes and descriptions are used to identify the nature of the failure. For detailed information on SIMs related to the global-active device environment, related pair status information indicating the volume with the most recent data, and recovery recommendations, see the *Global-Active Device User Guide* and the *Provisioning Guide* for your storage system. When the failure has been resolved, copy pairs need to be re-synchronized using CCI, the Replication tab, or Replication Manager.

HCS can be used to perform a limited number of recovery tasks that are included in this section.

Related tasks

- Recovering from HBA-related host path failures on page 364

Recovering from HBA-related host path failures

Remedy the HBA-related host path failure that is preventing host I/O to the storage system.

A SIM is not generated from a host path failure. However, there is a failure on the physical path between the host and the storage system.

Procedure

1. Using multi-path or SAN management software, identify the failure location.
2. Remedy the failure by doing one of the following:
   - Reconnect the physical path or reconfigure the SAN to recover from the path failure.
   - If you confirm a failed HBA, replace it.
3. Using multi-path software, resume I/O for the host on the recovered path.

Note that I/Os from the host might be resumed automatically.

Result

The failed path is operating again.

Related concepts

- About global-active device failures on page 363

Related tasks

- Editing LUN paths when exchanging a failed HBA on page 243

Discontinuing a global-active device environment

To discontinue the use of a global-active device environment, you must delete remote path settings, delete and unvirtualize the quorum disk, delete the command devices for pair management servers, and delete the virtual storage machine.

Before you begin

- Data on affected global-active device pairs has been saved.
- All global-active device pairs are unallocated, or the global-active device S-VOLs unallocated leaving basic volumes in use by hosts.

Procedure

1. Delete global-active device remote path settings.
   a. On the **Resources** tab, right-click the target storage system, and select **Remote Connections**.
   b. To delete a remote path, select a remote connection to be deleted in the **Connections (To)** tab, and then select **Remove Remote Paths**.
   c. To delete all remote paths allocated to the secondary storage system on the primary storage system, select the remote connections to be deleted in the **Connections (To)** tab, and then select **Remove Remote Connections**.
2. Delete and unvirtualize the global-active device quorum disk.
   a. On the **Resources** tab, in the **Storage Systems** tree, select and then right-click the target storage system. From the menu, select **Remote Connections**.
   b. In the **Remote Connections** window, on the **Quorum disk** tab, select the quorum disk, confirm it is correct, and delete it.

   Note: To delete a quorum disk, all global-active device pairs that use the quorum disk must be unallocated.

c. From **General Tasks**, select **Unvirtualize Volumes**. In the **Unvirtualize Volumes** dialog box, select the external storage system (quorum) and the internal storage system (primary or secondary storage system) and unvirtualize the external volume (quorum disk). Be sure to perform this for both the primary and secondary storage.

3. Delete the command devices for global-active device pair management servers.
   a. In the **Unallocate Volumes** dialog box, unallocate the pair management server and the command device.
   b. On the **Resources** tab, in the **Storage Systems** tree, select the primary or secondary storage system, and select **Volumes**. In the application area, click the **System GUI** link.
   c. From the volume list (LDEVs) that appear, select the command device that you want to delete.

   Tip: Look for `Command Device` in the **Attribute** column.

   d. From **More Actions**, select **Edit Command Devices** and in the dialog box that appears, delete the command device.
4. Delete the global-active device virtual storage machine and remove the secondary storage system resources from the global-active device virtual storage machine.

   Note: Before removing secondary storage system resources, you must unallocate the volumes on the global-active device virtual storage machine.

   a. On the **Administration** tab, select **Virtual Storage Machine**.
   b. Select the virtual storage machine related to global-active device pairs and click **Edit Virtual Storage Machine**.
   c. In the **Edit Virtual Storage Machine** window, select the secondary storage system, and then select **Remove Storage Systems**.

Result

The global-active device environment no longer exists.

Related tasks

- Unallocating global-active device pairs on page 360
- Unallocating individual global-active device volumes on page 362

9

Migrating data using nondisruptive migration (VSP G1000, VSP G1500, VSP F1500 and VSP Gx00 models)

The nondisruptive migration functionality allows the migration of host data (volumes and copy pair configurations) from an existing storage system to a newer storage system while maintaining continuous host I/O. Continuous host I/O is accomplished by using physical connections between storage systems for data migration, physical host I/O paths to the new storage system, and by using global storage virtualization features, such as virtual storage machines and virtual IDs assigned to volumes.

- ☐ About nondisruptive migration
- ☐ About migration planning
- ☐ Setting up migration projects and jobs
- ☐ Migrating volumes to the target storage system
- ☐ Migrating copy pair configurations to the target storage system
- ☐ Managing migration projects and jobs
- ☐ Post-migration tasks

About nondisruptive migration

Nondisruptive migration is a licensed HCS feature that makes it possible to relocate data from data from an existing (source) storage system to a new (target) storage system without interrupting access by hosts. This is accomplished using the global storage virtualization technology of Hitachi Virtual Storage Platform G1000 and G1500 (VSP G1000 and G1500), Hitachi Virtual Storage Platform G200, G400, G600, G800 (VSP Gx00 models), and Hitachi Virtual Storage Platform F1500 (VSP F1500). Resources on the source storage system are virtualized on the target storage system. From the perspective of the host, I/O requests continue to be serviced by the source storage system during the migration process.

The following storage system combinations are supported:

Source	Target
Hitachi Universal Storage Platform V/VM	VSP G1000, VSP G1500, and VSP F1500
Hitachi Universal Storage Platform V/VM	VSP Gx00 models
Hitachi Virtual Storage Platform	VSP G1000, VSP G1500, and VSP F1500
Hitachi Unified Storage VM	VSP G1000, VSP G1500, and VSP F1500
Hitachi Unified Storage VM	VSP Gx00 models

You can migrate the following types of data:

- Volumes
- Copy pair configurations (ShadowImage)

To migrate volumes and copy pair configurations, you create a migration project for a source and target storage system. You create one or more migration jobs within a migration project. Then, you perform individual migration steps for each migration job to complete the migration over time at a sustainable work load.

When you create a migration project, a virtual storage machine is created on the target storage system. This virtual storage machine contains the information about the source storage system. Source volumes that are selected for migration are virtualized, masking the physical characteristics of the target storage system, and the data is migrated to target volumes on the target storage system. As far as the host is concerned, I/O requests are still being serviced by the source target system during the migration process.

Note: In this guide, *target volume* means the virtualized source volume that resolves to the target volume, and *target volume for data migration* means the physical target storage system volume to which migration data will be written.

HCS nondisruptive migration has several benefits:

- Data is migrated between storage systems without interrupting host applications.
- You can maintain data replication throughout the migration process by allowing the target storage system to inherit pair configurations before migrating the actual data.
- You can reduce the overall migration effort by importing configuration definition files instead of having to reconfigure pairs on the target storage system.
- The migration process is designed to be carried out in stages to reduce demands on network bandwidth.
- You can easily monitor migration project and migration job progress and status by reviewing both numerical and graphical data, which includes estimated information about how long the migration is likely to take.
- Up to seven source storage systems can be consolidated into a single target VSP G1000, VSP G1500, VSP F1500 or VSP Gx00 models.

For detailed information about nondisruptive migration, including supported source and target storage systems and their required microcode or firmware versions, see the *Nondisruptive Migration User Guide* and *Hitachi Command Suite System Requirements*.

Related concepts

- About migration planning on page 369
- Understanding migration projects and jobs on page 370

About migration planning

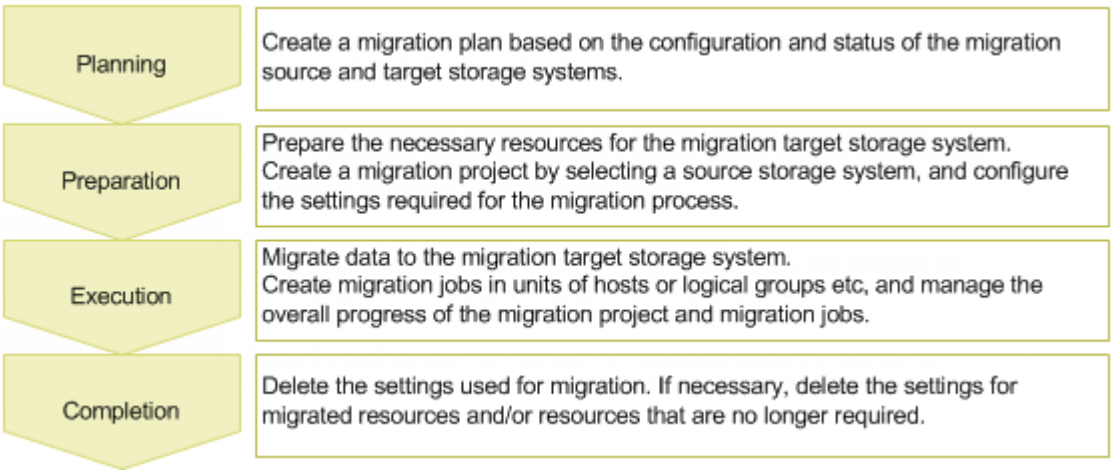
To successfully perform nondisruptive migration, detailed and careful planning is required before using HCS to implement nondisruptive migration.

Planning tasks include:

- Identifying the capacity of volumes and copy pairs to be migrated from a source storage system to a target storage system.
- Estimating migration overhead (when to migrate, how much data to migrate, and order of migration).
- Identifying target storage system ports to use as external ports for connecting the source and target storage systems for data migration.
- Identifying target storage system ports to be used as alternate paths for hosts whose volumes and copy pairs are being migrated from the source storage system.
- Identifying path management software, host OS software, or Fibre Channel switches to use to bring alternate paths online for hosts.

Caution: You must carefully plan a nondisruptive migration. For complete planning information, including concepts, requirements, limits, and restrictions, see the *Nondisruptive Migration User Guide*.

After you create a comprehensive migration plan, you can use HCS wizards for the preparation, execution, and completion phases of nondisruptive migration, as illustrated in the following workflow.



Related concepts

- About nondisruptive migration on page 368

Related tasks

- Preparing storage resources for migration on page 375

Setting up migration projects and jobs

Before migrating host volumes and copy pair configurations from a source storage system to a target storage system, you create migration projects, prepare storage resources, and then create migration jobs to manage the migration workflow.

Understanding migration projects and jobs

Migration projects and jobs define the scope of nondisruptive migration work to be performed for a source and target storage system. You can create migration projects and jobs (units of work) for up to seven source storage systems connected to one target storage system.

A *migration project* consists of a single target storage system paired with a single source storage system by using the Create Migration Project wizard. The migration project you create functions as a container for migration jobs.

A *migration job* consists of the volumes and copy pair configurations to be migrated to the target storage system. You select copy pair configurations,

hosts, logical groups, or manually selected volumes when you create a migration job. One or more migration jobs can be created for each migration project.

Because migrating a large number of volumes simultaneously is time-consuming and creates a heavy I/O workload, dividing a migration project into one or more migration jobs helps control the I/O load on both the source and target storage systems.

Migration projects and migration jobs appear on the Mobility tab, in the navigation pane under Migration Projects.

Note: If storage systems are not listed in Migration Projects, target storage systems are not yet registered. If, when you expand the tree for a registered target storage system, source storage systems are not listed, create a migration project for the target storage system and add a source storage system.

Related concepts

- Navigating migration projects and jobs on page 371

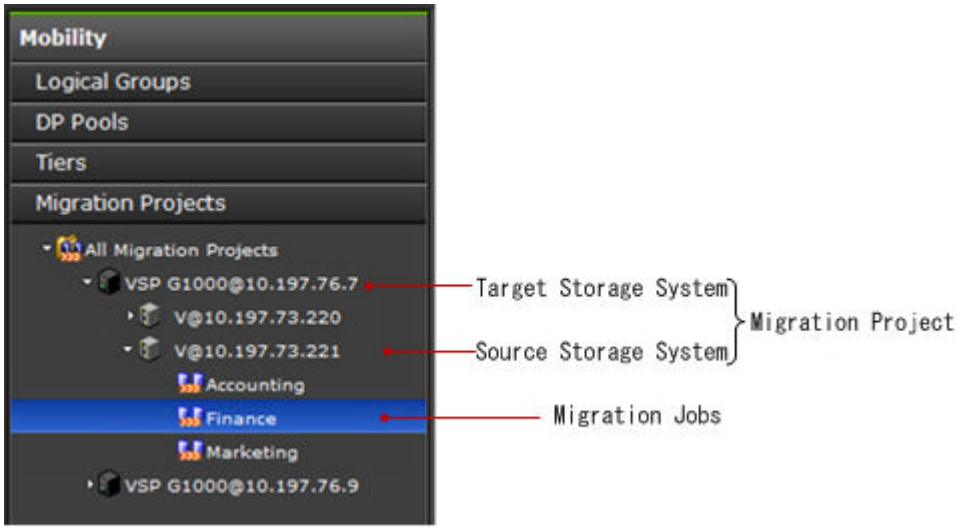
Related tasks

- Preparing storage resources for migration on page 375
- Creating a migration project on page 376
- Creating a migration job on page 378

Navigating migration projects and jobs

The migration user interface is flexible and provides various methods for performing migration tasks.

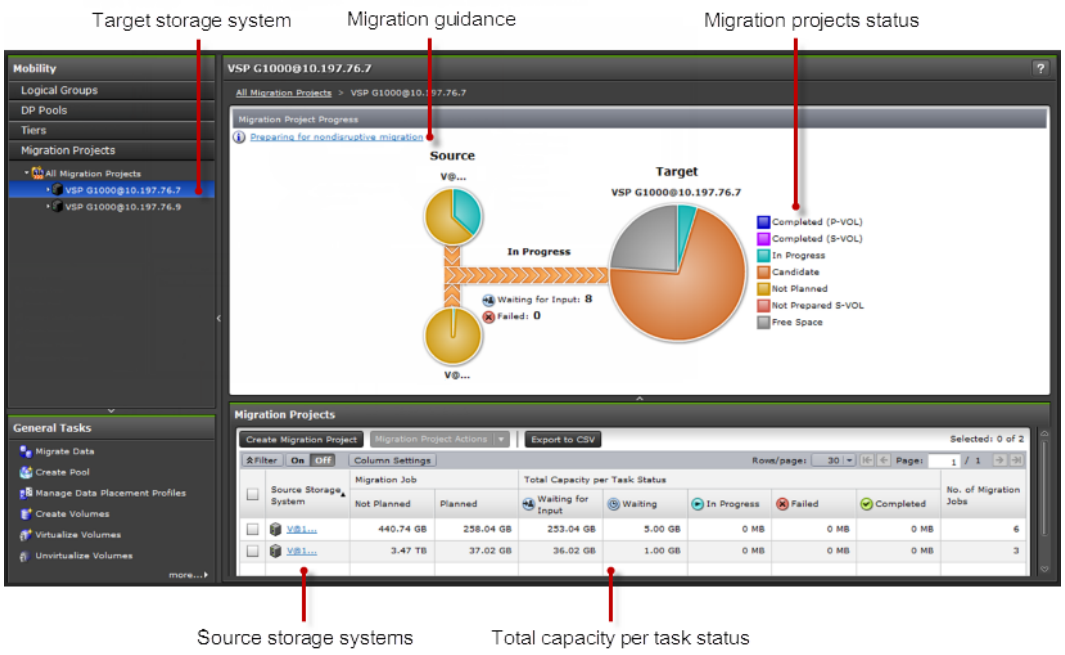
On the Mobility tab, when you select Migration Projects and then expand the tree, one or more target storage systems are listed.



- Each target storage system can be paired with up to seven (for consolidation purposes) nested source storage systems. This pairing comprises a migration project.
- For each source storage system in the tree, one or more nested migration jobs are listed.

When you select a target storage system from the navigation pane:

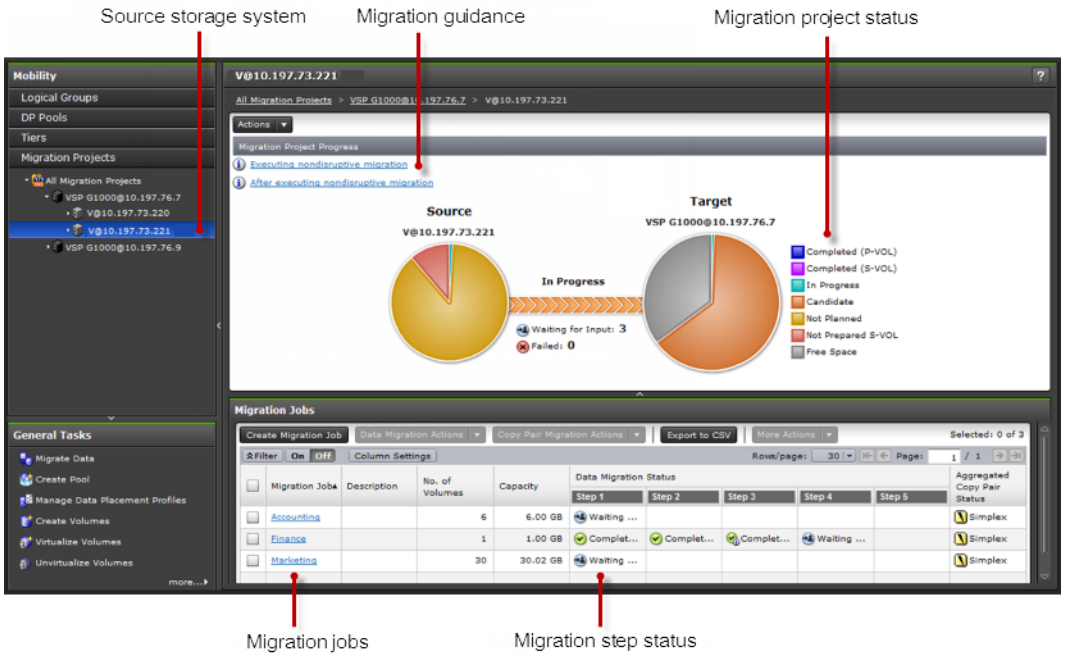
- Up to seven source storage systems (migration projects) are listed in the source storage systems list in the application pane. Migration project progress for all migration projects and for the entire target storage system is shown with color-coded pie charts labeled by status categories, which include corresponding volume capacity related to the status category. Using a mouse, hover over the pie-charts to see the status category name and related capacity metric, for example, the available free space on the target storage system, or the volume capacity that has been migrated.



- You click links in the Migration Project Progress pane to display migration guidance.
- In the Migration Projects pane, in the table of source storage systems, you can view the overall migration job status (Total Capacity per Task Status).

When you select a source storage system from the navigation pane:

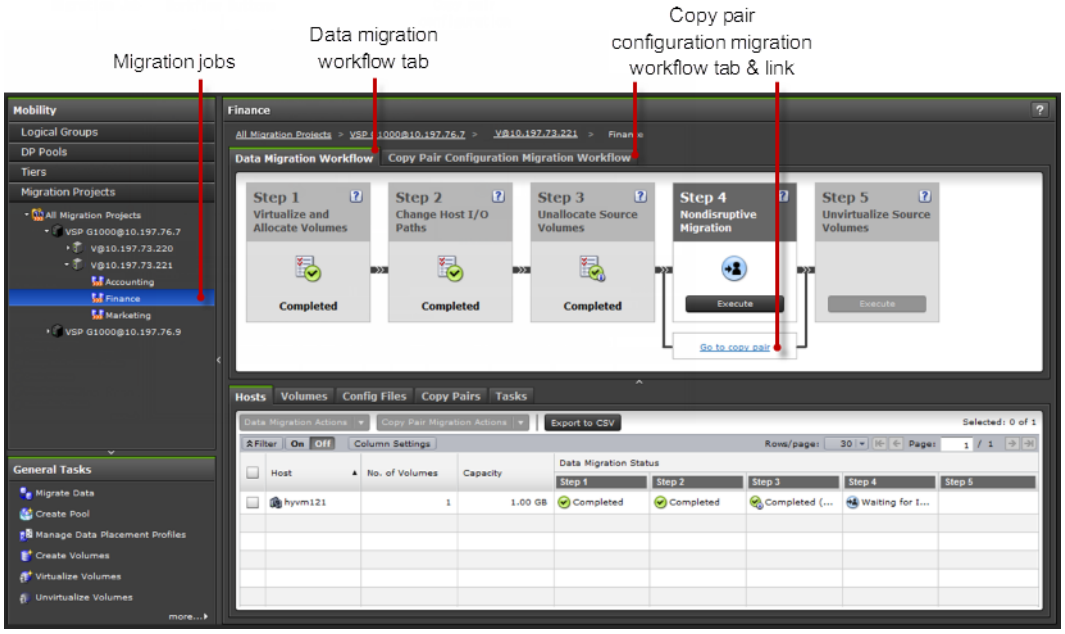
- In the Migration Project Progress pane, you can see the migration project progress and status for a single migration project (the selected source and target storage system pair).

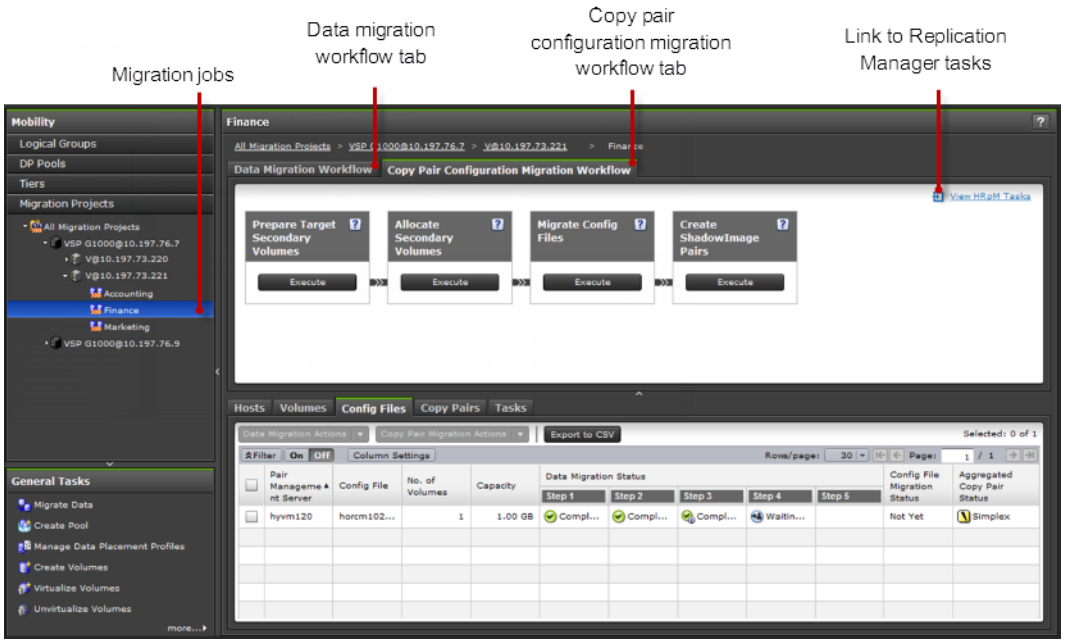


- You can click links in the Migration Project Progress pane to display guidance. Use the Actions menu to access migration project tasks.
- In the Migration Jobs section, you can create migration jobs. You can view or click links for all migration jobs for the migration project. You can view the progress and status of each migration job, including the volume count, total capacity to be migrated, and the status of each step in the migration job.

When you select a migration job from the navigation pane, or click the link for a migration job from the Migration Jobs list, GUI windows guide you through the steps in the workflow:

- Migration workflow tabs display for data migration and copy pair configuration steps. If you are migrating copy pair configurations as part of your migration project, a copy pair link is provided as a reminder of where the copy pair workflow is inserted into the data migration workflow. If a step is not complete, a button for performing the step is activated.





- In the bottom pane, tabs are available for viewing a migration job by Hosts, Volumes, Config Files (configuration definition files), and Copy Pairs (volumes belonging to a copy pair). Use Column Settings to configure the displayed information for each tab. The Tasks tab provides convenient access to job-specific migration tasks, and provides the Manually Change Task Status button to assist in recovering a stalled migration job.

Note: You can also click the Tasks & Alerts tab to review migration tasks.

Related concepts

- Understanding migration projects and jobs on page 370

Related tasks

- Tracking migration project status on page 393
- Tracking migration job status on page 394

Preparing storage resources for migration

Before you can begin the migration, you must prepare the storage resources and create a migration project.

Before you begin

- Install the nondisruptive migration license on the target storage system.
- Create a migration plan for each source and target storage system to be paired.
- Register the target storage system in HCS.

Note: For information about requirements and planning a migration, see the *Nondisruptive Migration User Guide*.

Procedure

1. On the **Mobility** tab, select **Migration Projects**.
2. Expand the tree and select a target storage system.
3. Select the target system under **All Migration Projects**. The **Preparation** screen includes steps that prompt you to perform the necessary tasks.
4. Perform the tasks described in Step 1, including adding storage systems and hosts, and (if migrating existing copy pairs) creating and allocating a command device (**Allocate Volumes**). This includes the following:
   - Register a host (pair management server) that manages the configuration definition files in HCS, and refresh the storage system.
   - When the target storage system is VSP G1000, VSP G1500, or VSP F1500, set the port for data migration.
   - Create the target storage system command device and allocate it to the pair management server.
   - Create the command device in the resource group (meta_resource) on the target storage system whose resource group ID is 0.
   - To connect the command device, use a port that is different from the one that manages the migration target volume that is used for nondisruptive migration.

- Disable user authentication for the command device.
- Change instances of HORCM_LDEVG to HORCM_LDEV. You cannot migrate configuration definition files in which HORCM_LDEVG is defined. The list of available configuration definition files will not include any configuration definition files in which HORCM_LDEVG is defined.

5. Start the Create Migration Project wizard (Step 2).
6. Create the necessary DP pools (Step 3).
7. Select the source storage system from the Migration Projects table at the bottom of the screen (Step 4) and proceed with the migration.

Related concepts

- Methods for registering hosts on page 66

Related tasks

- Creating a migration project on page 376
- Registering a storage system on page 60
- Installing the software on page 49
- Allocating volumes from general tasks on page 204
- Editing virtual storage machines on page 330
- Creating a DP pool on page 173
- Editing command devices on page 284

Creating a migration project

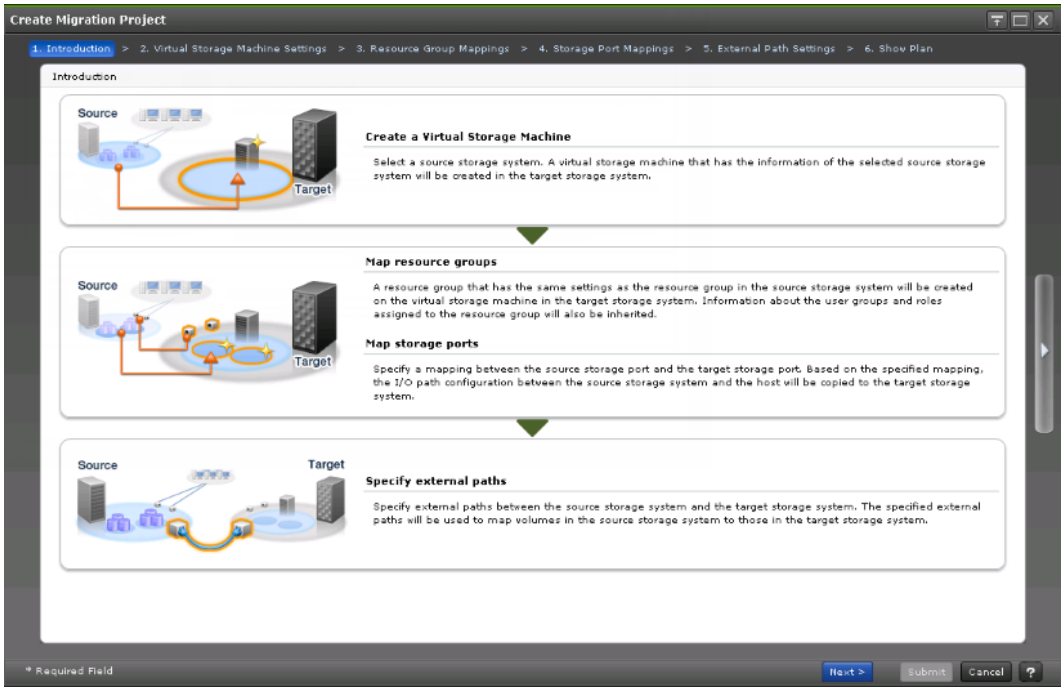
Use the Create Migration Project wizard to pair target and source storage systems, which make up a migration project.

Before you begin

- Review and meet the requirements for preparing storage resources, depending on your migration plan.

Procedure

1. On the **Mobility** tab, select **Migration Projects**.
2. Expand the tree and select a target storage system.
3. Click **Create Migration Project** to launch the Create Migration Project wizard.



4. Review the **Introduction** page, and then click **Next**.
5. On the **Virtual Storage Machine Settings** page:
   a. Select a source storage system.
   b. Accept the displayed virtual storage machine name, or enter a valid virtual storage machine name.
   c. (Optional) Select the **Migrate SI Pair Configuration** option to enable migration of copy pair configurations.
   If you select this option, do not make any changes to the configuration definition files until the migration is complete.

Note:

- Valid configuration definition files are indicated by **Yes** in the **Can the Copy Pair be Migrated** column. If **No** appears in this column for one or more of the configuration definition files you want to migrate, click **Cancel**, fix the configuration definition files, refresh hosts and the storage system, and then create the migration project again. For more information, see the planning chapter of the *Nondisruptive Migration User Guide*.
- When you finish creating the migration project, all source storage system configuration definition files are copied to the following locations, whether or not configuration definition files are valid:

  For Windows:

  Device-Manager-agent-installation-folder\agent\config\migration\backup

  For UNIX:

Device-Manager-agent-installation-directory/agent/config/migration/backup

6. Click **Next**.
7. On the **Resource Group Mappings** page, click **Next**.
8. On the **Storage Port Mappings** page, specify the mapping between the source and target storage ports, and then click **Next**.
   You can verify the port details for each selected target storage port in the **Port Details** table.
9. On the **External Path Settings** page, configure the external paths between the port of the source storage system and the externalor bidirectional port of the target storage system. Click **Increase Priority** or **Decrease Priority** to adjust the priority of the path. Click **Edit External Paths** to verify the external path configuration, to edit external paths using the topographical map or selection table, or to clear external path settings. When you are finished, click **Next**.
10. On the **Show Plan** page, confirm that the information in the plan summary and plan details are correct. If changes are required, click **Back**.
11. (Optional) Update the task name and provide a description.
12. (Optional) Expand **Schedule** to specify the task schedule.
    You can schedule the task to run immediately or later. The default setting is **Now**. If the task is scheduled to run immediately, you can select **View task status** to monitor the task after it is submitted.
13. Click **Submit.**
    If the task is scheduled to run immediately, the process begins.
14. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.

Result

The migration project is created, and the source storage system is displayed in the Migration Projects tree.

Related concepts

- Understanding migration projects and jobs on page 370

Related tasks

- Preparing storage resources for migration on page 375
- Deleting a migration project on page 400
- Creating a migration job on page 378

Creating a migration job

A migration job consists of source volumes (which might be copy pair P-VOLs) to be migrated to the target storage system. The volumes are selected by specifying one or more hosts or logical groups, by manually selecting

volumes, or by specifying configuration definition files containing copy pair information.

Before you begin

A migration project must exist for the source and target storage systems.

Procedure

1. On the **Mobility** tab, select **Migration Projects**.
2. Expand the tree and select a source storage system.
3. Click **Create Migration Job**.
4. Specify a name and description for the migration job.
5. Specify volumes by selecting one of the following options:

Option	Description
Config Files	Select **Config Files** when migrating copy pair configurations. This option is available only if you enabled migration of copy pair configurations when you created the migration project. If you select configuration definition files that define secondary volumes, also select the configuration definition files that define the corresponding primary volumes.
Hosts	Select **Hosts**, and then click **Select a Host** or **Select Hosts** to select multiple hosts.
Logical Group	Select **Logical Group**, and then click **Select a Logical Group**.
Manual	Select **Manual**, and then click **Add Volumes** in the Selected Volumes table to select source volumes for migration.

6. (Optional) Update the task name and provide a description.
7. (Optional) Select **View task status** to monitor the task after it is submitted.
8. Click **Submit**.

Result

In the migration projects list, the new migration job appears for the source storage system.

Related concepts

- Understanding migration projects and jobs on page 370
- Volume and copy pair configuration migration workflow on page 380

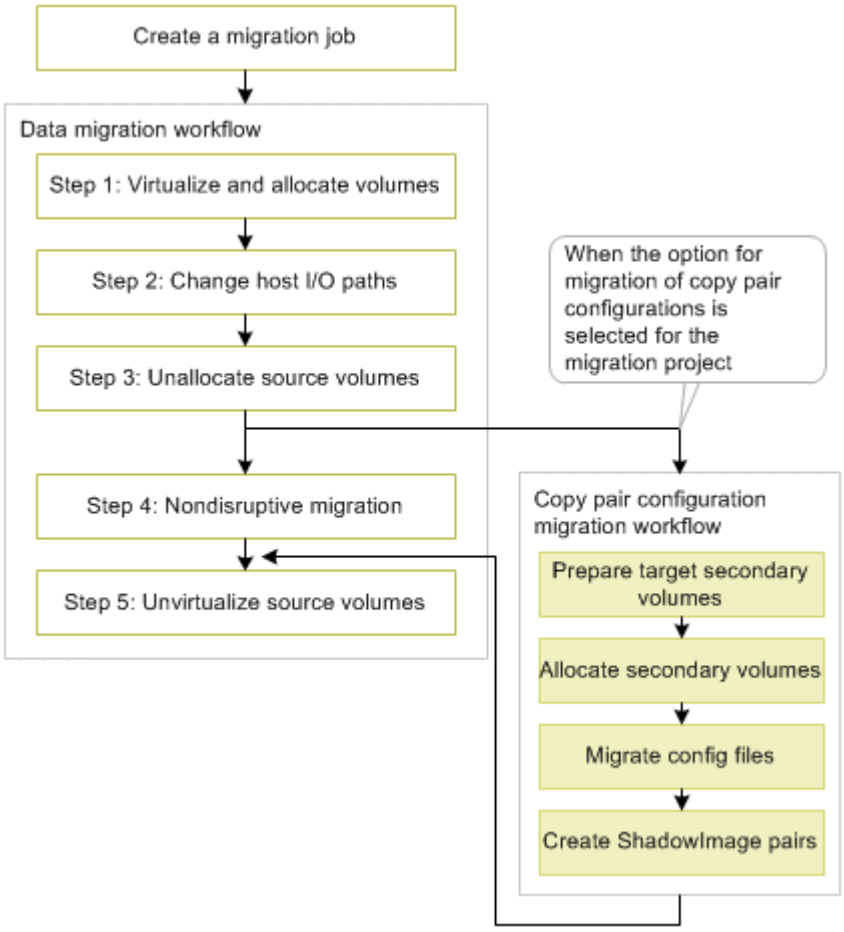
Migrating volumes to the target storage system

Volume and copy pair configuration migrations follow a standard workflow. The status of each step in the workflow is displayed as each step is performed so that you can track the progress of the migration. When all steps are completed for all volumes in the migration job, the migration job is complete. When all migration jobs are completed, the migration project is complete.

Volume and copy pair configuration migration workflow

HCS migration workflows support the migration of volumes and copy pair configurations.

The migration of individual volumes and copy pair P-VOLs is supported by the data migration workflow. Optional additional steps are required to prepare S-VOLs on the target storage system using the copy pair configuration migration workflow. The copy pair configuration workflow is performed as part of the data migration workflow.



- **Step 1: Virtualize and Allocate Volumes**
  Virtualize the volumes of the source storage system (including copy pair P-VOLs). Source volume IDs become target storage system virtual IDs. Source and target volume mappings and host group allocations are created.
- **Step 2: Change Host I/O Paths**
  To change paths, use path management software or change zoning settings. Bring host I/O paths online to the target storage system (for example, by Fibre Channel switch zoning), and take host I/O paths offline to the source storage system. These tasks require user confirmation that the host I/O paths have been changed.
- **Step 3: Unallocate Source Volumes**
  Delete the host LUN paths to the source volumes (including copy pair P-VOLs). The cache mode of the target volumes is automatically set to Write Sync so that source volumes continue to be updated before data migration.
  For migrating copy pair configurations, additional steps are required:

  Note: To maintain copy pair redundancy during volume migration (recommended), complete the copy pair configuration migration workflow before performing nondisruptive migration (step 4). Alternatively, you can perform the copy pair configuration migration workflow after nondisruptive migration (step 4) is complete for the P-VOL, and before unvirtualizing source volumes (step 5).

  - **Prepare Target Secondary Volumes**
    On the target storage system, create volumes to use as S-VOLs for P-VOLs whose data will be migrated. The source S-VOL ID is used as the virtual LDEV ID of target S-VOL. Assign target S-VOLs to a host group.
  - **Allocate Secondary Volumes**
    Delete the paths from S-VOLs on the source storage system to the backup server, and then allocate the S-VOLs on the target storage system to the backup server. Remove the backup server WWN from the source host group, and add the backup server WWN to the target host group.
  - **Migrate Config Files**
    Select source storage system configuration definition files to be migrated to the target storage system environment.
  - **Create ShadowImage Pairs**
    Reestablish source storage system copy pairs on the target storage system using the Change Pair Status Wizard in Replication Manager.
- **Step 4: Nondisruptive Migration**
  Migrate source volume data to the target volume for data migration.

Note: In this guide, *target volume* means the virtualized source volume that resolves to the target volume, and *target volume for data migration* means the physical target storage system volume to which migration data will be written.

- **Step 5: Unvirtualize Source Volumes**
  After data migration is completed, unvirtualize the source volumes. If the source volume is a copy pair volume, delete the copy pair from the source storage system before unvirtualizing the source volume.

Related tasks

- Creating a migration job on page 378
- Step 1: Virtualize and Allocate Volumes on page 382
- Step 2: Change host I/O paths on page 383
- Step 3: Unallocate source volumes on page 385
- Step 4: Nondisruptive migration on page 386
- Step 5: Unvirtualize source volumes on page 387
- Preparing target secondary volumes on page 388
- Allocating target secondary volumes on page 390
- Migrating configuration definition files on page 390
- Creating copy pairs on the target storage system on page 391

Step 1: Virtualize and Allocate Volumes

To migrate data, you must virtualize source volumes and allocate them to the target storage system.

Note: After you start the migration job (Step 1), we recommend that you complete the data migration without canceling it. (If you cancel an ongoing data migration to restore operations to the migration source storage system, you might need to manually execute the CLI.)

Before you begin

- A migration job must exist for the migration project.
- Verify that the target host is not registered in host groups that have different host modes or host mode options.
- Verify that if the target host is already registered in a host group of the target storage system, that the WWNs of the host group of the target storage system and corresponding host group of the source storage system are the same.

Procedure

1. On the **Mobility** tab, select **Migration Projects**.
2. Expand the tree for a source storage system, and select a migration job.

3. On the **Data Migration Workflow** tab, click **Execute** for **Step 1: Virtualize and Allocate Volumes**.

   Note: To resume a step for one or more migration job resources, for example, resources with a failed step task or sub-task, select the resources, and from the **Data Migration Actions** menu, select the step to resume.

4. On the **Virtualize Volumes Setting** page, verify the source and target storage systems and the migration job name.
5. (Optional) Change the parity group start number setting.

   For example, even though parity groups are automatically numbered starting from 1, you can add another source storage system to a target storage system and select a start number to easily differentiate the source of the parity groups.
6. (Optional) Change the cache logical partition (CLPR) to be used for a migration.
7. (Optional) Specify the initial number of the LDEV ID for the volumes to be added to the virtual storage machine.
8. Verify the external path information, and then click **Next**.
9. On the **Allocate Volumes Setting** page, click the **Volume Path** tab, and confirm the volume path mapping between source and target volumes.
10. Click the **Host Group** tab and click the provided links to verify host group mapping and other information, such as the scope of **LUN Security** and that the **Host Group Attributes** have not changed for the source and target.
11. Click **Edit Host Group Setting** to verify or change the host group name and the resource group for the host group.
12. Click **Next.**
13. On the **Show Plan** page, verify the plan summary and plan details, optionally add a task name and description, and then click **Submit**.
14. Confirm this task is **Completed** for the migration job.

Result

Source volumes are virtualized and allocated to hosts in the target storage system.

Related tasks

- Step 2: Change host I/O paths on page 383
- Preparing storage resources for migration on page 375

Step 2: Change host I/O paths

To redirect host I/O to the target storage system, host I/O paths must be online for the target storage system and offline for the source storage system.

Before you begin

Volumes must be virtualized and allocated.

Note: The following information applies to steps 4 and 6 of this procedure. The **Change Host I/O Paths** window only displays information about path changes. To change paths, you must use path management software or change zoning settings.

Note: If the LUN paths of source storage systems are left online and you proceed to the next step, data might be lost. For the target hosts, make sure that the LUN paths of all of the source storage systems are offline, and then submit the task.

Procedure

1. On the **Mobility** tab, select **Migration Projects**.
2. Expand the tree for a source storage system, and select a migration job.
3. On the **Data Migration Workflow** tab, click **Execute** for **Step 2: Change Host I/O Paths**.

   Note: To resume performing a step on one or more migration job resources, for example, resources with a failed step task or sub-task, select the resources, and from the **Data Migration Actions** menu, select the step to resume.

4. In the **Change Host I/O Paths** window, carefully review the **Step 1** instructions for bringing the host I/O path of the target volume online.

   To help you review and make the necessary changes, click **Export to CSV** to create a file that you can reference for the target LUN path information.
5. In **Step 2**, verify the target LUN path information, and then select the check box labeled **I have confirmed that all target LUN paths are online**.
6. Review the **Step 3** instructions for taking the host I/O path of the source volume offline. To help you review the source LUN path information, click **Export to CSV** to create a file that you can reference.
7. In **Step 4**, verify the source LUN path information, and then select the check box labeled **I have confirmed that all source LUN paths are offline**.
8. (Optional) Update the task name and provide a description.
9. Click **OK** to submit the task.
10. Confirm this task is **Completed** for the migration job.

Result

Host I/O paths are online to the target storage system and offline for the source storage system.

Related tasks

- Step 1: Virtualize and Allocate Volumes on page 382
- Step 3: Unallocate source volumes on page 385
- Preparing storage resources for migration on page 375

Step 3: Unallocate source volumes

Delete the source storage system LUN paths for hosts with migration source volumes, and change the cache mode on the target storage system.

Before you begin

Host I/O paths must be online for the target storage system and offline for the source storage system.

Procedure

1. On the **Mobility** tab, select **Migration Projects**.
2. Expand the tree for a source storage system, and select a migration job.
3. On the **Data Migration Workflow** tab, click **Execute** for **Step 3: Unallocate Source Volumes**.

   Note: To resume performing a step on one or more migration job resources, for example, resources with a failed step task or sub-task, select the resources, and from the **Data Migration Actions** menu, select the step to resume.

4. Verify the source LUN path information.

   The target storage system cache mode is automatically changed from Through to Write Sync to forward write I/O to the source storage system volume.
5. (Optional) Update the task name and provide a description.
6. (Optional) Expand **Schedule** to specify the task schedule.

   You can schedule the task to run immediately or later. The default setting is **Now**. If the task is scheduled to run immediately, you can select **View task status** to monitor the task after it is submitted.
7. Click **Submit.**
   If the task is scheduled to run immediately, the process begins.
8. Confirm this task is **Completed** for the migration job.

Result

Source volumes are unallocated from the host, the cache mode is set to Write Sync, and host I/O continues to update the source volume.

Next steps

If you are migrating copy pair configurations, follow the copy pair configuration migration workflow before continuing to Step 4: Nondisruptive Migration. Alternatively, you can perform the copy pair configuration migration workflow after nondisruptive migration (step 4) is complete for the P-VOL, but before you unvirtualize the source volumes (step 5).

Related tasks

- Step 2: Change host I/O paths on page 383
- Step 4: Nondisruptive migration on page 386
- Preparing target secondary volumes on page 388
- Preparing storage resources for migration on page 375

Step 4: Nondisruptive migration

Migrate source volume data to the target volume for data migration.

Note: To maintain copy pair redundancy during volume migration (recommended), complete the copy pair configuration migration workflow before performing nondisruptive migration (step 4). Alternatively, you can perform the copy pair configuration migration workflow after nondisruptive migration (step 4) is complete for the P-VOL, but before unvirtualizing source volumes (step 5).

Before you begin

Source volumes must be unallocated.

Procedure

1. On the **Mobility** tab, select **Migration Projects**.
2. Expand the tree for a source storage system, and select a migration job.
3. On the **Data Migration Workflow** tab, click **Execute** for **Step 4: Nondisruptive Migration**.

   Note: To resume performing a step on one or more migration job resources (such as resources with a failed step task or sub-task),select the resources, and from the **Data Migration Actions** menu, select the step to resume.

4. Verify the source and target storage systems and migration job name.
   The estimated migration time is shown for the migration job.
5. Click **Set Target Volumes for Data Migration** to specify a specific DP pool or parity group for the target volume for data migration, and then select either **Create new volumes** or **Use existing volumes**. Click **OK**.
6. (Optional) Select **Zero Page Reclaim** to reclaim capacity after migration. This option is enabled by default. Zero page reclaim can take

a long time. The estimated time is shown next to the check box. If you clear the check box, you can perform zero page reclaim after the migration job completes.

7. (Optional) Select the option to send email notification to a specified email address when the data migration completes.

   Note: Email notifications for Tiered Storage Manager events must be set up in advance. For more information see the *Hitachi Command Suite Administrator Guide.*

8. (Optional) Update the task name and provide a description.
9. (Optional) Expand **Schedule** to specify the task schedule.
   You can schedule the task to run immediately or later. The default setting is **Now**. If the task is scheduled to run immediately, you can select **View task status** to monitor the task after it is submitted.
10. Click **Submit.**
    If the task is scheduled to run immediately, the process begins.
11. Confirm this task is **Completed** for the migration job.

Result

Source volume data is migrated to target volumes for data migration.

Note: If the migration task status is "In progress" and the migration status of each volume remains at 99%, go to the Edit Local Replica Options window in Device Manager - Storage Navigator, disable the Nondisruptive Migration Data Consistency option, and then continue operation.

Related tasks

- Step 3: Unallocate source volumes on page 385
- Step 5: Unvirtualize source volumes on page 387
- Preparing storage resources for migration on page 375

Step 5: Unvirtualize source volumes

Unvirtualize source storage system volumes in the target storage system.

Before you begin

- The migration of the data must be completed.
- Copy pair configuration migration workflow steps (started after step 3 or step 4) must be completed.
- If the source volume is a ShadowImage copy pair volume, delete the source storage system copy pair before you unvirtualize the volume.

Procedure

1. On the **Mobility** tab, select **Migration Projects**.
2. Expand the tree for a source storage system, and select a migration job.

3. On the **Data Migration Workflow** tab, click **Execute** for **Step 5: Unvirtualize Source Volumes**.

   Note: To resume performing a step on one or more migration job resources, for example, resources with a failed step task or sub-task, select the resources, and from the **Data Migration Actions** menu, select the step to resume.

4. Verify that the Plan Summary information is correct.
5. Verify the list of source volumes to unvirtualize and their corresponding internal volume and parity group information.
6. Verify the external paths to be unallocated.
7. (Optional) Update the task name and provide a description.
8. (Optional) Expand **Schedule** to specify the task schedule.
   You can schedule the task to run immediately or later. The default setting is **Now**. If the task is scheduled to run immediately, you can select **View task status** to monitor the task after it is submitted.
9. Click **Submit.**
   If the task is scheduled to run immediately, the process begins.
10. Confirm this task is **Completed** for the migration job.

Result

Source volumes are unvirtualized and the migration job is completed.

Related concepts

- Volume and copy pair configuration migration workflow on page 380

Related tasks

- Step 4: Nondisruptive migration on page 386
- Preparing storage resources for migration on page 375

Migrating copy pair configurations to the target storage system

Copy pair configurations can be migrated from the source storage system to the target storage system using the copy pair configuration migration workflow. Copy pair redundancy can be maintained throughout the migration process. Configuration definition files are backed up, and are inherited by the target storage system for reestablishing copy pairs.

Preparing target secondary volumes

Create S-VOLs on the target storage system. The source S-VOL ID becomes the virtual LDEV ID of target storage system S-VOLs.

Before you begin

- Source volumes must be unallocated.
- A target storage system command device must be created and allocated to the pair management server.
- The option to migrate copy pair configurations was selected when the migration project was created.

Procedure

1. On the **Mobility** tab, select **Migration Projects**.
2. Expand the tree for a source storage system, and select a migration job.
3. Verify that steps 1 through step 3 on the **Data Migration Workflow** tab are completed, and then click the **Go to copy pair** link.
4. On the **Copy Pair Configuration Migration Workflow** tab, click **Execute** for **Prepare Target Secondary Volumes**.
5. On the **Create Target Secondary Volumes** page, select the S-VOL, and then click **Set Target Secondary Volumes** to specify a specific DP Pool or Parity Group for the target S-VOL, and then select **Create new volumes** or **Use existing volumes**. If necessary, specify the initial number of the LDEV ID for the volumes to be created on the virtual storage machine, and then click **OK**.
6. In the **Target Pools/Parity Groups Information** area, review the capacity changes for DP pools or parity groups.
7. Repeat steps 5 and 6 for each displayed source and target S-VOL pair, and then click **Next**.
8. On the **Set Paths for Target Secondary Volumes** page, compare the LUN path between the source secondary volume and the backup server with the corresponding LUN path setting of the target storage system. On the **Volumes** tab, you can see which backup server is allocated to each volume. On the **Host Groups** tab, you can compare the host group settings. As required, edit the names of the host groups created on the target storage system and the resource groups in which the host groups will be created. Enter the necessary settings, and then click **Next**.
9. Review the plan, and then click **Submit**.
10. Confirm this task is completed.

Result

Target storage system secondary volumes are prepared.

Related tasks

- Step 3: Unallocate source volumes on page 385
- Allocating target secondary volumes on page 390
- Preparing storage resources for migration on page 375

Allocating target secondary volumes

Delete the backup server paths for source storage system S-VOLs, and then establish backup server paths for target storage system S-VOLs.

Before you begin

- Source volumes must be unallocated.
- The target storage system secondary volumes are prepared.
- Stop the backup script that is running on the backup server.

Procedure

1. On the **Mobility** tab, select **Migration Projects**.
2. Expand the tree for a source storage system, and select a migration job.
3. On the **Copy Pair Configuration Migration Workflow** tab, click **Execute** for **Allocate Secondary Volumes**.
4. Note the host group information for the source and target storage systems, which will assist you when you are changing the backup server paths from the source storage system to the target storage system.
5. Delete the source storage system backup server WWN:
   a. Launch **Element Manager** to run the native management tool for the source storage system to delete the backup server WWN from the host group of the source storage system.
   b. Verify that the backup server WWN has been deleted from the host group of the source storage system.
6. Add the backup server WWN to the target storage system:
   a. On the **Resources** tab, for the target storage system, select **Ports/ Host Groups/iSCSI Targets** and add the backup server WWN to the host group of the target storage system.
   b. Verify that the backup server WWN has been added to the host group of the target storage system.

Result

Backup server paths for source storage system S-VOLs are deleted, and backup server paths for target storage system S-VOLs are created.

Related tasks

- Step 3: Unallocate source volumes on page 385
- Migrating configuration definition files on page 390
- Preparing storage resources for migration on page 375
- Starting Element Manager on page 575

Migrating configuration definition files

For the migration job, selected configuration definition files can be migrated for use on the target storage system.

Before you begin

- Source volumes must be unallocated.
- Target secondary volumes must be allocated.

Procedure

1. On the **Mobility** tab, select **Migration Projects**.
2. Expand the tree for a source storage system, and select a migration job.
3. On the **Copy Pair Configuration Migration Workflow** tab, click **Migrate Config Files**.
4. Click **Add/Remove Config Files**.

   Note: The configuration definition file to be migrated must be specified in HORCM_LDEV format. Files in which HORCM_LDEVG is defined cannot be migrated and does not appear in the list.

5. Click **Add** and **Remove** to populate the list of configuration definition files to be migrated, and then click **OK**.
6. (Optional) Update the task name and provide a description.
7. (Optional) Expand **Schedule** to specify the task schedule.

   You can schedule the task to run immediately or later. The default setting is **Now**. If the task is scheduled to run immediately, you can select **View task status** to monitor the task after it is submitted.
8. Click **Submit.**
   If the task is scheduled to run immediately, the process begins.
9. Confirm this task is **Completed** for the migration job.

Result

Configuration definition files are migrated for use on the target storage system. HORCM instances for the migrated configuration definition files are restarted. If a HORCM instance is running on the source storage system, the configuration definition file is migrated to the target storage system.

Related tasks

- Step 3: Unallocate source volumes on page 385
- Creating copy pairs on the target storage system on page 391
- Preparing storage resources for migration on page 375

Creating copy pairs on the target storage system

Create ShadowImage pairs on the target storage system based on migrated configuration definition files.

Before you begin

- Source volumes must be unallocated.
- Migration of configuration definition files must be completed.

- Review the limitations for creating copy pairs in the *Nondisruptive Migration User Guide*.

Procedure

1. On the **Mobility** tab, select **Migration Projects**.
2. Expand the tree for a source storage system, and select a migration job.
3. On the **Copy Pair Configuration Migration Workflow** tab, click **Execute** for **Create ShadowImage Pairs**.
4. In the **Create ShadowImage Pairs** dialog box, select the copy group you want.
5. Click **Create Copy Pairs**.
6. In the Replication Manager **Change Pair Status Wizard**, on the **Select Copy Pairs** page, select the copy pair you want and click **Next**.
7. On the **Select Pair Operation** page, review the copy pair information and click **Next**.
8. On the **Set Schedule** page, select **Execute on Schedule** to schedule the task for a specific date and time, or select **Execute Immediately**, and then click **Next**.

   Note: You cannot edit the created task. If you want to reschedule the task, delete it and create a new task.

9. On the **Confirm** page, verify the copy pair settings, select the check box labeled **Yes. I have confirmed the above information and wish to change pair status**, and then click **Confirm**.
10. On the **Finish** page, confirm the task and then click **Finish**. The task is registered as a Replication Manager task.
11. To verify the status of the Replication Manager task:
    a. On the **Copy Pair Configuration Migration Workflow** tab, click **View HRpM Tasks** to display the Replication Manager task list.
    b. Open the **Create ShadowImage Pairs** dialog box and check **Status**.

Result

Copy pairs are created on the target storage system.

Next steps

Restart the backup script of the backup server for the target storage system if necessary.

Related tasks

- Step 4: Nondisruptive migration on page 386
- Preparing storage resources for migration on page 375

Managing migration projects and jobs

After you create migration projects and jobs, you can monitor them and perform maintenance tasks on them.

Tracking migration project status

The progress of migration projects is tracked and displayed graphically.

The following aggregated status information from migration jobs is displayed:

- The total capacity of volumes for a migration project.
  - You can estimate the time required to complete the migration project by reviewing the remaining capacity and up-to-date elapsed time for migrations.
  - You can determine if you need to add physical storage capacity, or create additional DP pools, by comparing the data capacity to be migrated against the free capacity of the target storage system.
- The completion state of migration jobs, which provides a snapshot of remaining migration activity for project completion.

Before you begin

A migration project with migration jobs must exist.

Procedure

1. On the **Mobility** tab, select **Migration Projects**.
2. Expand the tree and select a target storage system.
3. Track project status. In the **Migration Project Progress** window, categories (by capacity) are displayed in a pie chart format (using a variety of colors) for both source and target storage systems. Use your mouse to hover over the pie chart slices to review the name of the category and related statistics. The following table describes the color-coded pie chart categories and other status information:

Pie chart category	Storage system	Description
Completed	Source and Target	The total capacity of the volumes whose data migration is complete, meaning that the (nondisruptive migration) step of the data migration workflow is complete.
Completed (P-VOL)	Source and Target	If you are migrating copy pair configurations, the total capacity of the source storage system primary volumes whose data migration to the target storage system is complete.
Completed (S-VOL)	Target	The total capacity of volumes that are created for secondary volumes in the target storage system.

Pie chart category	Storage system	Description
In Progress	Source and Target	The total capacity of the source storage system volumes that are registered in a migration job and whose data migration to the target storage system is not yet completed. Tasks in this category include those that are in the **Waiting for Action** (waiting for user input to continue the migration job) or **Failed** (waiting for recovery of a stalled migration job) status.
Candidate	Target	The total capacity of the volumes that are in the source storage system and are not registered in a migration job. If you migrate copy pair configurations, the total capacity of the volumes for secondary volumes created in the target storage system is included.
Not Planned	Source	The logical capacity of the volumes that are in the source storage system and are not registered in a migration job.
Not Prepared S-VOL	Source	The total capacity of the secondary volumes of the copy pairs in the source storage system minus the total capacity of secondary volumes in the target storage system.
Free Space	Target	The total value of the free capacity in the target storage system minus the total capacity of the volumes in the In Progress or Candidate categories in the source storage system.

If free space warnings display, you might need to adjust target storage resources that are available for migration. Migration jobs that have the Waiting for Action status need to be continued, and stalled (Failed status) migration jobs need to be investigated and resumed. For information about recovery or troubleshooting tasks, see the *Nondisruptive Migration User Guide*.

Related tasks

- Tracking migration job status on page 394

Tracking migration job status

The status of migration jobs (and steps within the jobs) is displayed for each migration job.

Before you begin

Migration jobs must exist within a migration project.

Procedure

1. On the **Mobility** tab, select **Migration Projects**.
2. Expand the tree and select a source storage system.
3. In the **Migration Jobs** list, review the **Data Migration Status** column for one or more listed migration jobs.

   Data migration status values are:

Status	Description
Waiting for Action	Waiting for you to take action.
Waiting	Scheduled task.
In Progress	Running task.
Completed	Successful task.
Failed	Unsuccessful step or substep in the task. You must determine the cause and perform recovery steps. For detailed troubleshooting information to recover from conditions that caused a Failed status at the migration job step or substep level, see the *Nondisruptive Migration User Guide.*
Waiting for Action (Manually recover)	Waiting for you to take action to recover from a previously failed step or substep in a task. This status indicates that the recovery process involves reverting to a previous state, before the failed task occurred. You set this status as part of the recovery process, and then resume the task to complete the step.
Completed (Manually recover)	You have resolved an issue that caused a failed state. You set this status as part of the recovery process to clear the failure and allow the migration to proceed.

4. To review detailed migration job status, select a migration job from the tree, or click a migration job link.
5. Check the status of migration steps by clicking the **Hosts**, **Volumes**, **Config Files**, or **Copy Pairs** tabs.

Tip: The Volumes tab provides links for Source and Target volumes, and for each step status, providing direct access to volume details and task status information. Column Settings is configurable for relating volumes to a variety of resources, such as host or virtualized information used for migration.

6. Click the **Tasks** tab to see all migration job related tasks, or click the **Tasks & Alerts** tab, and then click **HCS Tasks** to see all migration tasks for all migration jobs.

Note: To recover from conditions that cause a **Failed** status at the migration job step or substep level, see the *Nondisruptive Migration User Guide* for detailed troubleshooting information and recommendations.

Related tasks

- Tracking migration project status on page 393
- Changing the status of a failed task on page 396

About resuming failed migration jobs

A migration job is performed in a series of steps, each step consisting of one task and one or more sub-tasks. If the task or one or more sub-tasks fails, a status of Failed is displayed for the step.

You cannot continue a migration job until you do one of the following items:

- You must correct the failed task or failed sub-tasks for the step, and then change status of the failed migration job so that the job can be resumed. You do this by:
  - Correcting the state of a resource, and then manually changing the failed job status to Completed (Manually recover) to indicate that the issue is corrected, which allows you to resume the job beginning with the next step.
  - Restoring a resource to its previous state before the task or sub-task was executed, and then manually changing the failed job status to Waiting for Action (Manually recover) to indicate that the issue is corrected, which allows you to resume the job at the failed step and then continue with remaining steps.
- Under some circumstances one or more migration job volumes might need to be deleted, and added to a new migration job.
- Under some circumstances, a migration job might need to be deleted and re-created.

For information about how to recover from conditions causing a failed status at the migration job step or sub-step level, see the *Nondisruptive Migration User Guide*, which provides detailed troubleshooting information and recommendations for specific steps and sub-steps, including guidance on restoring a resource to its previous state, or correcting the state of a resource through remedial actions.

Related tasks

- Tracking migration job status on page 394
- Changing the status of a failed task on page 396
- Canceling a migration job on page 399
- Canceling a volume from a migration job on page 400

Changing the status of a failed task

If a task fails, after recovery actions have been performed, change the status of the failed task to continue the migration job.

Before you begin

Correct the issue that is causing the failed task or subtask.

Caution: If you do not correct the issue that is causing a failed task or sub-task for a migration step before you change the status of the failed step to

Waiting for Action (Manually recover) or Completed (Manually recover) in the Manually Change Task Status dialog box, the step cannot be resumed by using the GUI.

You must cancel (delete) affected volumes from the migration job, return them to the state before the task ran, and then create a new migration job for the volumes.

For information about how to recover from conditions that cause a failed status at the migration job step or sub-step level, see the *Nondisruptive Migration User Guide*, which provides detailed troubleshooting information and recommendations for specific steps and sub-steps, including guidance on restoring a resource to its previous state, or correcting the state of a resource through remedial actions.

Procedure

1. On the **Mobility** tab, select **Migration Projects**.
2. Expand the tree, and select the migration job with the failed task.
3. On the **Volumes** or **Tasks** tab, do one of the following:
   - On the **Volumes** tab, select the volume for which recovery has been performed, and then from the **More Actions** menu, select **Manually Change Task Status**.
   - On the **Volumes** or **Tasks** tab, locate the task for which recovery has been performed, click the **Failed** status link, review task details, and click **Manually Change Task Status**.
4. In the **Manually Change Task Status** dialog box, verify the displayed information, and then select the **Status Change** check box to confirm that recovery has been performed.
5. From the **Select a Status** list, select one of the following:
   - **Waiting for Action (Manually recover)**: Resumes the task from the step during which the task failed.
   - **Completed (Manually recover)**: Resumes the task from the next step.
6. (Optional) Update the task name and provide a description.
7. (Optional) Select **View task status** to monitor the task after it is submitted.
8. Click **Submit**, and verify the status change for the migration step.

Result

The migration job step status is changed, and you can resume the migration job.

Related tasks

- Tracking migration job status on page 394
- Restarting a failed or stopped HCS task on page 510

Editing a migration project

After a migration project is created, and before performing data migration, you can edit migration project information.

You can edit the following migration project information:

- Settings for resource groups that are created in the virtual storage machine of the target storage system
- Settings that define the mapping between source and target storage ports
- Settings that define the external paths between the ports of the source storage system and the external or bidirectional ports of the target storage system

Procedure

1. On the **Mobility** tab, select **Migration Projects**.
2. Expand the tree and select a target storage system.
3. In the **Migration Projects** list, select a migration project.
4. From the **Migration Project Actions** menu, choose one of the following options:

Option	Description
Edit Resource Group Mapping	Select the resource groups on the target storage system to map to the resource groups on the source storage system, and then click **OK**. You can verify the new resource group mapping by opening the **Edit Resource Group Mapping** dialog box again. **Tip**: If you change the mapping to the resource groups on the target storage system, the information about the user groups and roles that are assigned to the resource groups on the source storage system are not inherited. Edit the resource group mapping, and then assign resource groups and roles to the user groups on the target storage system.
Edit Port Mapping	Select the ports of the target storage system to map to the ports of the source storage system, or change the destination of ports that are already mapped, and then click **OK**. You can verify the new I/O path mapping by opening the **Edit Port Mapping** dialog box again.
Edit External Path Setting	Specify paths in accordance with the displayed instructions, and then click Submit. The submitted plan is registered as a task. Check the task results in the Tasks & Alerts tab. You can verify the new external path

Option	Description
	mapping by opening the **Edit External Path Setting** dialog box again.

Result

Resource group mapping, port mapping, and external path settings are updated for the migration project.

Related concepts

- Understanding migration projects and jobs on page 370

Related tasks

- Preparing storage resources for migration on page 375
- Creating a migration project on page 376

Editing a migration job

You can change the name or description of a migration job.

Tip: If you want to migrate more volumes, create a new migration job. If you want to stop volume migration, you can cancel a migration job, or delete (cancel) volumes from a migration job.

Procedure

1. On the **Mobility** tab, select **Migration Projects**.
2. Expand the tree for a source storage system, and select a migration job.
3. From the **More Actions** menu, select **Edit Migration Job**.
4. Change the migration job name or description and click **OK**.
5. Verify the results by reviewing the migration job name and description in the **Migration Jobs** list.

Result

The migration job is updated.

Related tasks

- Creating a migration job on page 378

Canceling a migration job

You can cancel a migration job that has not been started, meaning no migration activity has occurred. After you cancel a migration job, its volumes can be specified as migration source volumes in another migration job.

You can also cancel a migration job in progress. However, to use the deleted volumes of migration jobs for other migration jobs, you must manually recover the volumes of the cancelled migration jobs. To do so, follow the

instructions in displayed messages. For details about how to manually recover the volume status, see the *Nondisruptive Migration User Guide*.

Procedure

1. On the **Mobility** tab, select **Migration Projects**.
2. Select a source storage system.
3. From the **Migration Jobs** list, select one or more migration jobs.
4. From the **More Actions** menu, select **Cancel Migration Jobs**.
5. Verify the displayed information, select the check box to confirm that you have reviewed and understand the plan, and then click **Submit**.
6. Confirm that the job is cancelled by reviewing the migration jobs list again.

Result

The migration job is cancelled.

Related tasks

- Creating a migration job on page 378

Canceling a volume from a migration job

You can delete individual volumes from a migration job, unless migration for the volume is already complete. Deleted volumes can be specified as migration source volumes in another migration job.

Procedure

1. On the **Mobility** tab, select **Migration Projects**.
2. Expand the tree for a source storage system, and select a migration job.
3. On the **Volumes** tab, select the volumes to delete.
4. From the **More Actions** menu, choose **Cancel Migration Volumes**.
5. Verify the displayed information, select the check box to confirm you have reviewed and understand the plan, and then click **Submit**.
6. Confirm the volume deletions by selecting the migration job again.

Result

The selected volumes are deleted from the migration job.

Related tasks

- Creating a migration job on page 378
- Editing a migration job on page 399

Deleting a migration project

You can delete a migration project that has no active tasks. You might want to delete a migration project after the migration is complete or if the

migration fails and you want to try again by creating another migration project. This procedure deletes the settings associated with the migration project. It does not delete the individual tasks in the project.

Before you begin

Verify that there are no tasks that are active for the migration jobs in the migration project.

Procedure

1. On the **Mobility** tab, select **Migration Projects**.
2. Expand the tree and select a target storage system.
3. In the **Migration Projects** list, select a migration project.
4. From the **Migration Project Actions** menu, select **Delete Migration Project**.
5. Verify that the plan is correct, and then click **Submit**.
   The plan is registered as a task.
6. Check the task results in the **Tasks & Alerts** tab.

Result

The migration project is deleted.

Related concepts

- Understanding migration projects and jobs on page 370

Related tasks

- Creating a migration project on page 376

Post-migration tasks

After a migration project is complete, you can perform clean up and storage optimization tasks. Review the tasks to determine if they are applicable to your completed migration project.

Shredding or deleting source volumes

You can use Hitachi Command Suite to shred or delete source storage system volumes, including volumes that are configured as copy pairs at the migration source, whose data migration is complete.

Before you begin

- Ensure that source storage system volumes or source storage system copy pairs to be shredded or deleted are unallocated.

Procedure

1. Shred the source storage system volumes and source storage system copy pair volumes (P-VOLs and S-VOLS) whose migration is complete.

   Note: Shredding writes bit patterns to a volume that overwrite the existing volume data. Original volume data cannot be restored after it is shredded.

2. Delete the source volumes and source storage system copy pair volumes (P-VOLs and S-VOLS) whose migration is complete.

   Note: You can delete the primary volume (source P-VOL) and secondary volumes (S-VOLS) at the same time.

Related concepts

- About shredding volume data on page 155

Related tasks

- Deleting unallocated volumes on page 157

Removing the source storage system from Hitachi Command Suite

After a migration project is complete, and you have completed post-migration tasks such as volume shredding and deletion, you can remove (unregister) the source storage system from HCS.

You might want to remove a source storage system from HCS, for example, because the storage system is being decommissioned and physically removed, or it is being redeployed.

Related tasks

- Removing a storage system on page 63

Reclaiming zero pages for target volumes

You can use Hitachi Command Suite to reclaim zero pages for target volumes that were created using DP pools. You can reclaim zero pages if migration jobs are complete, but the zero page reclaim option was not specified.

Zero page reclaim looks for allocated but unused DP pool pages, and reclaims the empty pages for later allocation to DP volumes that require expansion.

Related tasks

- Reclaiming zero pages on page 186

Optimizing the target storage configuration

Migration adds host I/O load to the target storage system. You can tune the performance of the target storage system and configure settings that were not automatically migrated.

Performance tuning tasks can include:

- Setting the cache of the target storage system
- Tuning MP blade performance for the newly added host I/O loads
- Adding I/O paths for the host
- Configuring tier profiles for HDT volumes

Configuration setting tasks include:

- Configuring copy pairs and logical groups
- Configuring volume information, such as filter, label, and search values

Related concepts

- About optimizing HBA configurations on page 435

Related tasks

- Analyzing and load balancing MP Blades or MP Units on page 532
- Creating a CLPR on page 439
- Optimizing data placement in a DP pool on page 449

10

Configuring a VVol environment

This chapter describes the configuration and maintenance tasks in Hitachi Command Suite in order to support VMware VVols. There are also references to the virtual machine administrator (VM administrator) and the associated tasks for that function. Those tasks are specific to using vSphere and are used only when describing interrelated functions that are performed by storage administrators using HCS.

- Working with VMware VVols
- Storage Policy Based Management (SPBM)
- VVols and block storage
- VVols and file storage

Working with VMware VVols

In Hitachi Command Suite (HCS), virtual machine management includes support for VMware VVols.

When working with block storage, you can specify and group the DP pools (or Thin Image pools) to be disclosed to the Storage Provider for VMware vCenter (the VASA provider, also referred to as the virtual machine administrator), and then define the group of pools as a storage container. However, if you are working with VVols for file storage, then you need to specify the file systems that are associated with the storage container.

Additionally, you can set a capability profile, which groups metadata that are provided by the storage container (the datastore), such as IOPS, Latency, and Availability. Afterward, you can disclose the capability schema, default profiles (for block storage only), and storage container to the VM administrator.

In a general perspective, the tasks of creating virtual machines based on a policy and allocating volumes comprise Storage Policy Based Management (SPBM).

In HCS, VMware VVol management is mostly performed on the Administration tab. On the Administration tab, you can complete the following tasks:

- Create storage containers
- Adjust the capacity of storage containers
- Manage the capabilities associated with storage containers
- Specify the VASA provider when working with VVols for file storage
- Manage the components used by a virtualization server to access VVols (the protocol endpoints)
- Change capability profiles when VVol compliance issues are identified
- Display information about SLUs (VVols) on the storage container to check the relationship between virtual machines and SLUs.
- Delete storage containers

Note: Three components are required to implement a VMware VVol environment: A separate HCS system that is used only for VVol management (referred to in this manual as HCS for VVols), a VASA provider, and vSphere. This chapter focuses on the VVol related tasks that are performed using HCS. Before you start the configuration and management tasks that are described in this chapter, installation and setup tasks relating to those components should already have been completed. For details, see the *Hitachi Storage Provider for VMware vCenter Deployment Guide* and the *Hitachi Command Suite Administrator Guide*. If you manage storage system resources by using

Device Manager - Storage Navigator (instead of using HCS), you cannot launch Device Manager - Storage Navigator from HCS for VVols.

Related concepts

- About registering a host on page 66

Related tasks

- Registering a storage system on page 60

Storage Policy Based Management (SPBM)

Storage Policy Based Management (SPBM) comprises the tasks of creating virtual machines based on a policy and allocating volumes.

About Storage Policy Based Management

One of the prerequisites for configuring a VMware VVol environment involves defining the Storage Policy Based Management (SPBM). An SPBM is a structure that is based on a defined VM storage policy, where the VASA provider and Hitachi Command Suite (HCS) automatically select the storage resources that meet the VM storage policy, and creates VVols.

After DP pools (or Thin Image pools) have been specified and grouped (or, file systems have been specified when working with VVols for file storage), they can be defined as a storage container, for which a capability profile can be set. A capability profile defines the capabilities of a storage container (for example, IOPS, Latency, Availability). Afterward, both the storage container and associated capability profile can be disclosed to the VM administrator.

Note: After you change the configuration of a DP pool for block storage, or a file system for file storage, check and if needed, change the capability profile.

The following steps outline the general task workflow for creating an SPBM structure. The tasks that apply to the VMware vSphere administrator (VM administrator) are noted in parenthesis:

1. **Define capability profiles**. Set the capability profile in the storage container for each DP pool for block storage and for each file system, for file storage.
2. **Define VM Storage Policies** (applies to the VM administrator). Create virtual machine storage policies based on the contents of the capability schema. The content that is defined in the capability schema can be viewed as the data service from vCenters.
3. **Check storage policy compatibility** (applies to the VM administrator). The vCenter reports whether storage containers exist that meet the policy criteria. This is done by comparing the capability profiles and VM storage policies.
4. **Create a VVol** (applies to the VM administrator).

Related concepts

- About default profiles on page 408
- SPBM implementation examples on page 408

About default profiles

Depending on your environment, default profiles may be used.

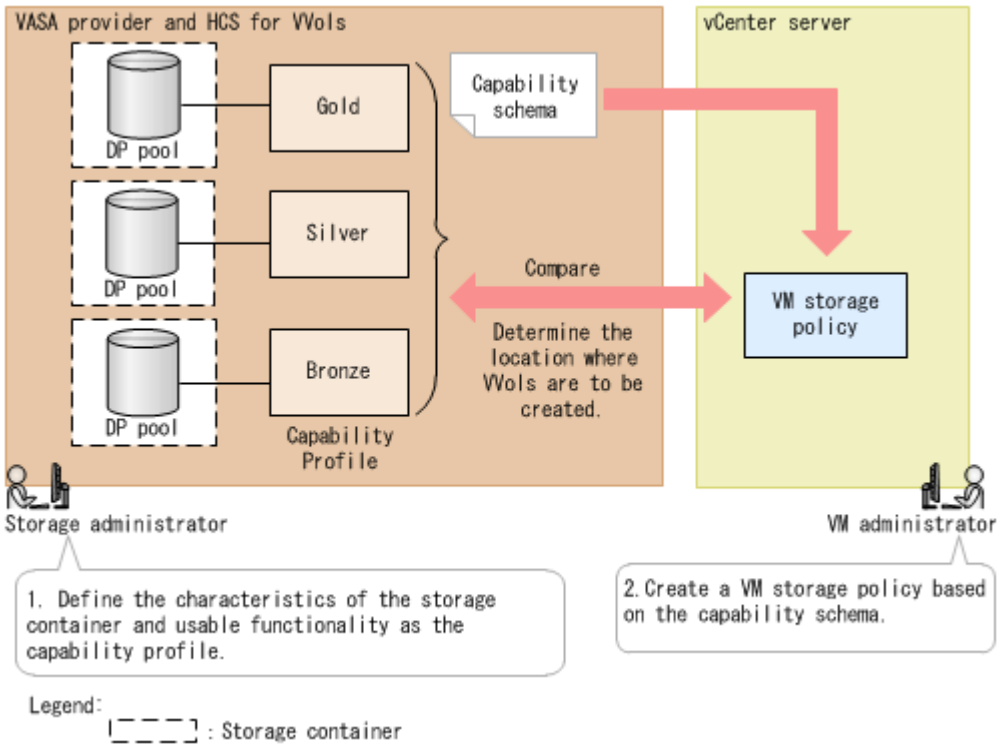
If the VM administrator creates virtual machines without defining a VM storage policy, or if virtual machines cannot be created, then default profiles are applied. The default profiles must be defined in advance for each storage container, as needed. By design, default profiles are specified so that every profile matches. If any profile does not match the default profile, then the default profile is specified. The default profile can be changed as needed.

Note: Default profiles only apply to VVols for block storage.

SPBM implementation examples

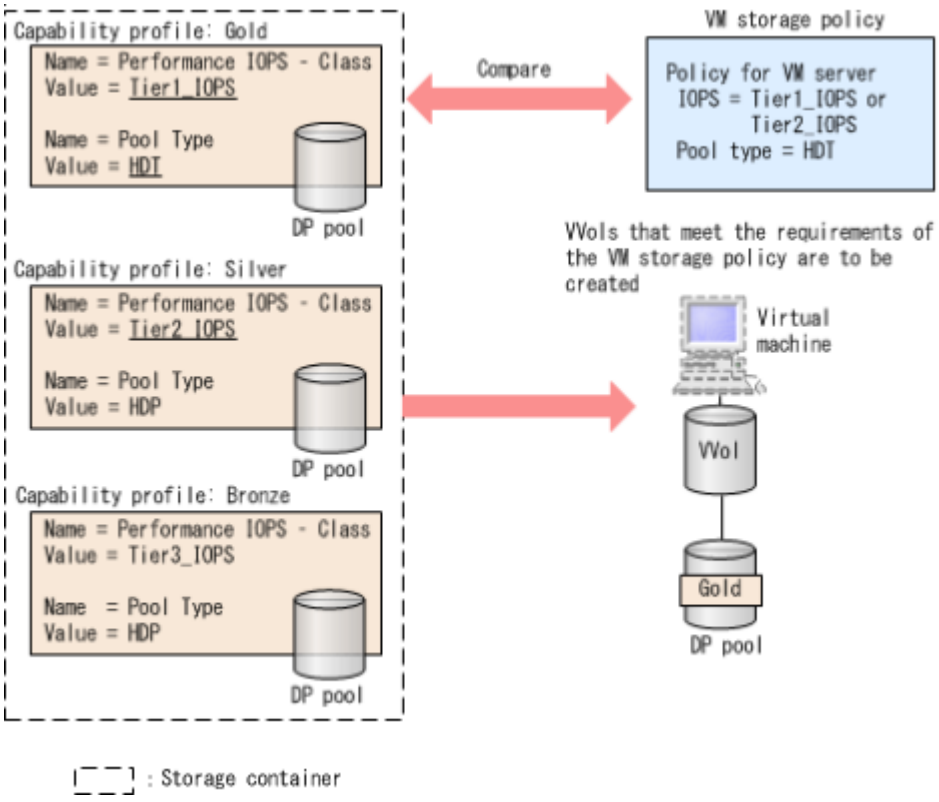
The workflow for establishing a VVol environment and the roles between the VASA provider and Hitachi Command Suite for VVols, are represented in the following SPBM implementation examples.

The following diagram provides an example of a typical SPBM implementation.



The following example provides details of the capability profile information that is associated with the storage container of the previous example. In the figure, the storage administrator has created the three types of storage capability profiles: Gold, Silver, and Bronze. The VM administrator has defined, as requirements for the virtual machine, IOPS as Tier1 or Tier2, and the pool type as HDT.

If the storage capability profile and the VM storage policy are compared, the storage capability profile that satisfies the requirements indicated by the VM administrator is Gold. In this case, the virtual machine is created from the DP pools that correspond to the Gold storage capability profile.



VVols and block storage

You can use block storage VVols in your environment after completing some initial configuration tasks. Afterward, you can manage your VVol environment for block storage by expanding the storage container capacity, deleting storage containers or removing the VVol environment entirely.

VVol configuration and operation tasks for block storage

The overall configuration and operation of a VMware VVol environment using Hitachi Command Suite involves a workflow that encompasses four main tasks:

- Preparing the environment

- Configuring the environment
- Managing the environment
- Removing the environment

Both HCS for VVols and HCS are required to complete the tasks.

Preparing the environment

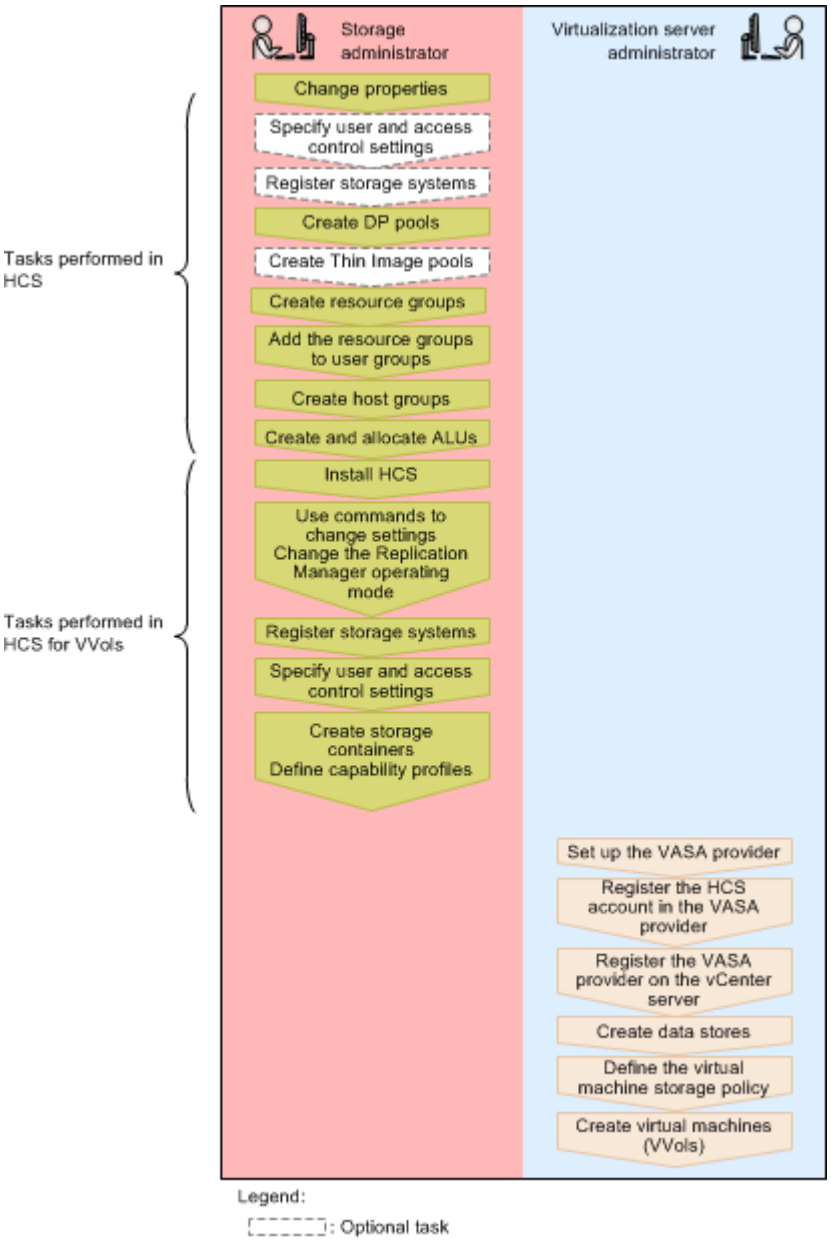
- After the VM administrator defines the requirements for the virtual machines, you can determine the necessary resources (such as DP pools, Thin Image pools, storage ports, LDEV IDs, and host group numbers) that will be needed, based on the requirements defined for the virtual machines.

- You can determine the configuration of resource groups according to the operations. We recommend registering resources for virtual machines separately, and to a different resource group for management.

Configuring the environment

- Changing HCS settings in order to work with and use VVols
- Specifying settings so that the VASA provider can access HCS for VVols. This includes the creation of a new user for the VASA provider, for multi-tenant operations. In the settings of the user group to which the user for the VASA provider belongs, assign an HCS role to the resource group for storage containers: Admin or Modify. Both the Admin role and Modify role are for All Resources.
- Using HCS to create DP pools and Thin Image pools, and registering them to the resource groups for storage containers
- Creating protocol endpoints (ALUs) and allocating them to virtualization servers
- Using HCS for VVols to define the Storage Policy Based Management (SPBM)
- Reviewing configuration guidelines (for example, configuring resource groups as required)

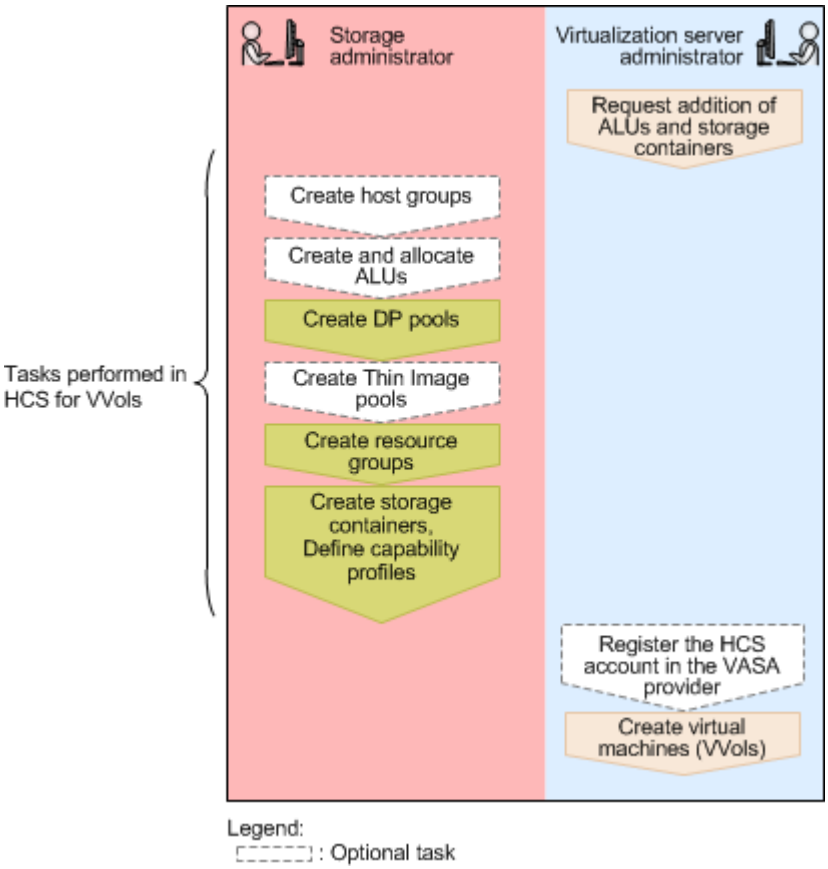
The following illustration provides an overview of these tasks.



Managing the environment

- Avoiding storage container capacity excesses
- Modifying capability profiles when VVol compliance is violated (monitoring virtual machines)
- Creating additional storage containers (snapshot and cloning activities)

After your site configures a VVol environment, the VM administrator might request the creation of ALUs and storage containers. The following workflow illustrates the tasks that are associated with that work.



Note: To create snapshots for block storage, create Thin Image pools and add them to a resource group for a storage container in advance. To create clones for block storage, add DP pools (and when creating clones of VMs that are powered on, Thin Image pools) for clones to a resource group for a storage container, and then create a storage container. To manage tasks after the environment that uses VVols is configured, use HCS for VVols.

Removing the environment

Deleting an obsolete VMware VVol environment

Licensing

For HCS for VVols, only register the license for Device Manager. If you are using HCS for operations, the HCS license that was previously registered, applies.

Related concepts

- Reviewing supported functions for VVol management on page 413
- Guidelines for working with virtual machines in a VMware VVol environment on page 414
- VVol resource group configurations on page 414
- Removing a VVol environment for block storage VVols on page 424

Related tasks

- Reviewing the tasks for creating block storage VVols on page 416
- Defining a protocol endpoint for block storage VVols on page 417
- Determining ALU and SLU volume attributes for block storage VVols on page 418
- Creating storage containers for block storage VVols on page 418
- Expanding the capacity of a storage container for block storage VVols on page 420
- Changing a capability profile for block storage VVols on page 422
- Deleting storage containers for block storage VVols on page 423
- Deleting a protocol endpoint for block storage VVols on page 425

Reviewing supported functions for VVol management

Some functions of other instances of Hitachi Command Suite (HCS) can be used for the resources that support VMware VVol management for block storage.

The following list of functions are supported in HCS for VVols, for the resources that are used for VMware VVol management (ALU, SLU, and DP pools including SLUs):

- Create, edit, and delete resource groups
- Edit, expand, and shrink DP pools
- Reclaim zero pages for DP volumes
- Manually monitor and relocate HDT pools
- Create, edit, and delete schedule templates for HDT pool monitoring and data relocation
- Enable or disable data relocation for HDT volumes
- Apply a tiering policy to HDT volumes
- Change the tiering policy definition for HDT volumes (customize a tiering policy)
- Edit the external LDEV tier rank for external HDT pool volumes
- Create, edit, and delete volume labels
- Search resources
- Analyze storage system performance

Some functions cannot be analyzed, and functions that are not included in the list above are not supported.

Note: Do not use DP pools in which ALUs and SLUs exist, for data migration. If these resources are used for data migration, the data might become corrupted. To migrate resources that have the SLU attribute, use VMware vSphere. Also, do not select ALU, SLU and DP pools (including SLUs) when manually selecting volumes in the following cases: Allocating volumes, allocating like volumes, creating volumes, deleting hosts, and creating file systems. If an ALU or SLU is selected, data might become corrupted.

Guidelines for working with virtual machines in a VMware VVol environment

In an environment that uses VMware VVol for block storage, the status of resources and the storage system might affect virtual machine operations. Following are the guidelines that your site should review before working with virtual machines in the VMware VVol environment.

- Confirm that the DP pools in the storage container have sufficient capacity. If there is insufficient capacity, then it might not be possible to create, start, suspend, or run clones for a virtual machine.
- Confirm that the Thin Image pools in the storage container have enough capacity. When the capacity for a Thin Image pool becomes limited, it might not be possible to create a snapshot or a clone of a virtual machine. If it is not possible, a new Thin Image pair might not be created, and you might have to delete all of the snapshot data. To expand the Thin Image pool capacity, do not create a snapshot until the expansion processing is complete.
- Specify a threshold value for a Thin Image pool. When the capacity for a Thin Image pool exceeds the threshold value, it might not be possible to create a snapshot or a clone of a virtual machine. For details about setting the threshold value for Thin Image, see the *Hitachi Thin Image User Guide*.
- After deleting a virtual machine or snapshots, the deletion of SLUs (VVols) might take a long time to complete. The capacity of the deleted SLU might not be applied to the DP pool and Thin Image pool immediately after deleting a virtual machine or snapshot.
- When the storage system configuration is changed in an environment other than Hitachi Command Suite (for example, using Device Manager), update the storage system information in HCS. If the storage system information in HCS for VVol is not up-to-date, then it might not be possible to create the virtual machine.
- When a virtual machine is operating, do not delete resource groups that contain SLUs. Also, do not delete a DP pool from which SLUs have been created and do not delete the LDEV IDs that are used for SLUs from a resource group or move them to another group. Otherwise, the virtual machines to which the relevant VVols are assigned might become inoperable.
- If you use more than one instance of HCS for VVols, do not share storage containers.

VVol resource group configurations

When resources for a storage system are divided and used by several storage administrators in a VMware VVol environment for block storage, you can use a specific resource group configuration.

Because a resource group that corresponds to a storage container is also used by the virtual machine (VM) administrator, only register the LDEV IDs to be used when using VVols. Register other resources to other resource groups.

To avoid having the administrator of the entire storage system add resources to each resource group, add extra LDEV IDs to a resource group that corresponds to a storage container, in advance. This will allow your site to continue operations while SLUs, DP pool volumes, or Thin Image pool volumes are being added, or while the capacity of the storage container is being expanded.

Resource Type	Item to add	Administrator Type	Target Resource Group
DP pool	DP pool	Storage administrator VM administrator	Storage container
Thin Image pool volume	LDEV ID	Storage administrator VM administrator	Storage container
Pool volume (additional)	LDEV ID	Storage administrator VM administrator	Storage container
SLU[1, 2, 3]	LDEV ID	Storage administrator VM administrator	Storage container
ALU	LDEV ID	Storage administrator VM administrator	Any
	Port number	Storage administrator	
	Host group or iSCSI target number	Storage administrator	
DP pool volume Thin Image pool volume	Parity group	Storage administrator	

Notes:

1. Only add an unused LDEV ID.
2. When registering LDEV IDs, register an LDEV ID that has a different SSID than an LDEV ID that has already been registered in other instances of HCS for VVols.
3. For storage systems whose resources are managed by a virtual storage machine, only add LDEV IDs for which virtual management settings are valid. For details on virtual management settings, see the *System Administrator Guide* (for VSP G1000, VSP G1500, and VSP F1500) or the *System Administrator Guide* (for VSP Gx00 models or VSP Fx00 models).

An SLU (VVol) is also required for tasks that are completed using virtual machines. The following table shows the required number of LDEV IDs when using VMware VVols for each purpose. Refer to this information when registering DP pools and LDEV IDs to the storage container, and when planning the DP pool capacity and LDEV IDs. In some cases, more LDEV IDs are required than those that are indicated in the following table. For this reason, make sure you have more LDEV IDs than indicated by the provided estimates in the table. For example, if a virtual machine that has two hard

disks and for which one memory-snapshot was obtained is running, the required number of LDEV IDs is 5.

Purpose	Required number of LDEV IDs (estimated)
Configuration management	1
Data	Same as the number of hard disks on the virtual machine
Activating a virtual machine	1
Obtaining snapshots of a virtual machine	Same as the number of memory-snapshots to be obtained for the virtual machine
Suspending a virtual machine	1

- **Managing resources in one storage container**. To use VVols in one storage container, register LDEV IDs to be used for a resource group that corresponds to that storage container. Register other resources to other resource groups for management.
- **Managing resources in multiple storage containers by sharing ALUs**. To use multiple storage containers, register ALUs as resources to be shared among the storage containers, to resource groups. Ensure that the resource groups are different from the resource groups that correspond to the storage containers for management.

Reviewing the tasks for creating block storage VVols

After you have completed Hitachi Command Suite (HCS) configuration tasks, you can use HCS for VVols to create storage containers and define capability profiles for them.

There are two main tasks for creating block storage VVols:

- **Configuring HCS**. After setting up a VASA provider, change settings within HCS so that VVols can be used. Afterward, DP pools and Thin Image pools can be created and registered to the resource groups for storage containers. Protocol endpoints (ALUs) must be created so they can be allocated to virtualization servers.

Note: An SLU (VVol) is a volume that stores data for virtual machines. When a virtual machine is created, an SLU is created from a DP pool, and then it is bound to an ALU through path allocation.

- **Creating VVols using HCS for VVols**. Install HCS for VVol, and then make any necessary changes to the settings. Create storage containers and define capability profiles for them.

Note: The installation, registration, and configuration of an HCS system dedicated to a VVol environment are required. When you register a storage system to an HCS system that is dedicated to a VVol, clear the Enable SVP/CCI user account authentication by HCS check box. Also, an account for

the storage system must be created in HCS for VVols in advance. You need to specify an account that is different from the one that was specified when the storage system was registered in HCS. The account to be registered must have a built-in user group specified for 'Administrator'. For VSP G400, G600, G800 or VSP Fx00 models with NAS modules where the microcode version is earlier than 83-04-0X-XX/XX, the account to be registered also must have a built-in user group specified for 'Support Personnel'. You can create the user account and password from Device Manager - Storage Navigator. When using more than one instance of HCS for VVols, you need to specify a different account for each instance.

For details about how to install HCS and how to change HCS settings to use VVols, see the *Hitachi Command Suite Installation and Configuration Guide* and the *Hitachi Command Suite Administrator Guide*.

For details about how to create an account for the storage system from Device Manager - Storage Navigator, see the *System Administrator Guide* for your storage system.

For details about how to install and set up a VASA provider, see the manuals for the VASA provider.

Defining a protocol endpoint for block storage VVols

To define a protocol endpoint, host groups must exist so that they can be allocated to an ALU. Then, the ALU can be allocated to a virtualization server.

Use the Set Protocol Endpoints dialog box to create host groups and allocate ALUs.

Note: When you create or delete protocol endpoints, the configuration information of the storage system changes. To continue tasks using HCS for VVols after the change, you may need to refresh the storage system.

Procedure

1. On the **Administration** tab, navigate to **Virtual Disks Management**, and select **VMware**.
2. On the **Configuration (Block Storage)** tab, select the target storage system, then click **Set Protocol Endpoints**.
3. On the **Set Protocol Endpoints** dialog box, click **Create Host Groups** to launch the **Ports/Host Groups/iSCSI Targets** window.
   a. On the **Ports/Host Groups/iSCSI Targets** window, click **Create Host Groups**, and establish host groups. Specify the following settings: **Host mode**: 21 (VMware Extension), **Host mode option**: 63 (VAAI) Support option for vStorage APIs that are based on T10 standards.
   b. To confirm whether the task has completed, click the **Check task status** link of the **Set Protocol Endpoints** dialog box to launch the

Tasks window. Proceed to step 4 only after confirming that the host group was created.

4. Click **Allocate ALUs** to launch the **Logical Device** window.
5. On the **Logical Device** window, click **Create LDEVs** to create an ALU and add a LUN path.

Determining ALU and SLU volume attributes for block storage VVols

Depending on the resources that are used for VMware VVol management (ALU, SLU, and DP pools, including SLUs), you can complete certain tasks, such as searching for resources. However, because there are some usage restrictions when working with certain resources (such as ALUs and SLUs), it is important to check whether volume attributes are ALU or SLU.

Before you begin

Because SLUs are not included in the capacity or volumes of the virtualization server, check the capacity and the volumes that are allocated to the virtualization server:

1. On the Resources tab, from Hosts, select the target virtualization server. In the list of volumes, check the displayed ALUs.
2. On the Resources tab, select the target storage system from Storage systems.
3. In the list of volumes displayed for Open-Reserved under Volumes, check the capacity and the volume of an SLU that is bound to an ALU.

Procedure

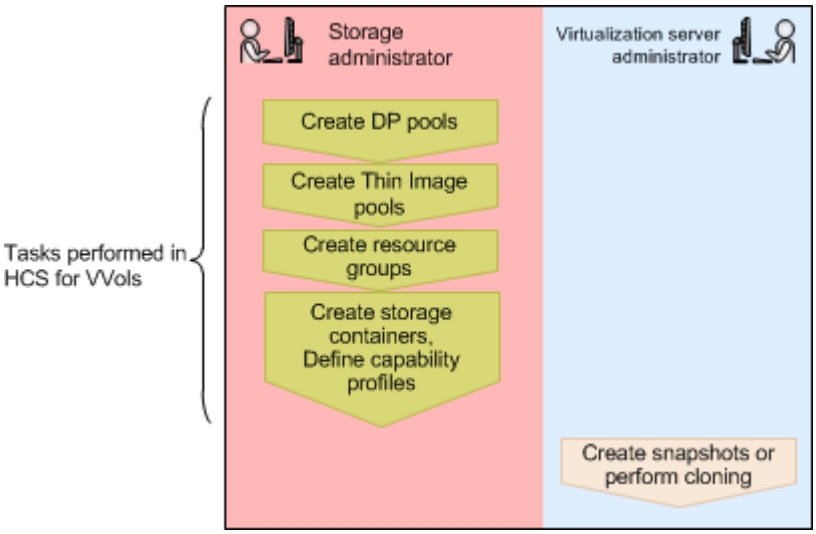
1. On the **Resources** tab, select the target storage system from **Storage systems**.
2. Expand the tree for **Volumes**.
   a. To check whether the volume attribute is **ALU**, select **Open-Allocated** or **Open-Unallocated**.
   b. To check whether the volume attribute is **SLU**, select **Open-Reserved**.
3. In the list of displayed volumes, check the volume attribute.
4. To check the LDEV ID of an ALU to which an SLU is bound, in the list of volumes displayed for **Open-Reserved**, click **Column Settings** to display **Binding ALU**.

Creating storage containers for block storage VVols

One of the tasks for defining the Storage Policy Based Management (SPBM) is creating a storage container.

You create a storage container so that it can be used by the virtualization server. The VM administrator can use the created storage container when creating a datastore.

To create snapshots and clones for a virtual machine using VMware VVols, you need to create storage containers. The following figure illustrates the workflow.



Before you begin

- A protocol endpoint (ALU) must be defined.
- To use storage containers for multiple virtualization servers, any user accounts that are used for a VASA Provider must be registered in advance, for each virtualization server.
- A resource group, DP pool, and if needed, a Thin Image (TI) pool for the storage container must be prepared. If a TI pool is created, then add the LDEV ID of the TI pool volume to the resource group.
- When you create or delete protocol endpoints, or edit, delete, create, or expand a TI pool, the configuration information of the storage system changes. To continue tasks using HCS for VVols after the change, you may need to refresh the storage system.

Note: In HCS, only TI pool volumes can be viewed. To add a TI pool volume to a resource group, note the TI pool ID when creating a TI pool. Open the Create/edit Resource Group dialog box, and then select the LDEV IDs tab. From the SS/TI Pool ID column, select a volume where the target TI pool ID is displayed, and then add it to the resource group. Afterward, all of the LDEV IDs for other volumes in the same TI pool are automatically added. Later on, if you need to expand, delete, or create other TI pools, launch Device Manager - Storage Navigator from the Resources tab, then right-click DP Pools for the target storage system and choose System GUI. To ensure that resource groups can be created, refresh the storage system.

Procedure

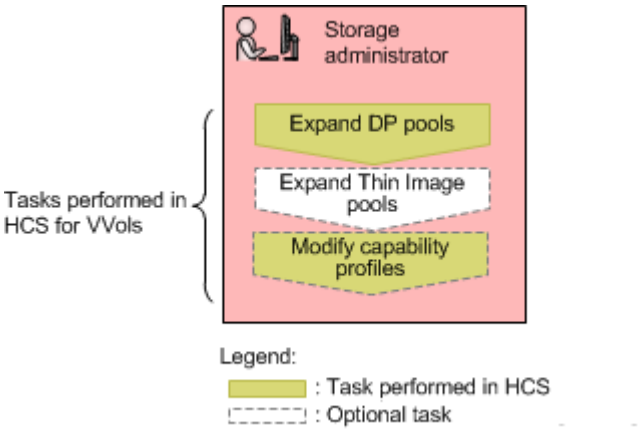
1. On the **Administration** tab, select **Virtual Disks Management**, click the **Vendors** tab, and then select **VMware**.
2. On the **VASA Providers** tab, click **VASA Provider (Block Storage)**.
3. On the **Storage Containers and SPBM** tab, click **Create Storage Container**.
4. Specify the following Basic Storage Container information:
   a. In the **Name** field, enter the name of the storage container.
   b. (Optional) In the **Description** field, enter a description.
   c. From the **Storage System** list, select the target storage system.
   d. From the **Resource Group** list, select the resource group that was used to register the pool for creating VVols, snapshots, and clones.
5. Define a capability profile for DP pools that are listed in the table for which the **Name** column indicates **Undefined**:
   a. Click **Define Profile** to create a new profile.
   b. In the **Name** field, enter the name of the capability profile.
   c. (Optional) In the **Description** field, enter a description.
   d. Within the **Managed Storage Capabilities** and **Auto-generated Capabilities** sections, either specify settings or check existing settings to ensure that the storage container to be created will match the capabilities of the VM storage policy.
   e. Click **OK** to save the new capability profile and return to the **Create Storage Container** dialog box.
6. (Optional) Default profiles are hidden. To view and modify default profiles, click **Advanced Options**.
7. Click **Submit**.
   The completed operation will be registered as a task.

Next steps

Check the results of the completed task on the Tasks & Alerts tab. Then, after the task has completed, you can check for the created storage container from the Storage Containers list.

Expanding the capacity of a storage container for block storage VVols

If the capacity of a storage container is about to exceed the threshold, then you can expand the pool capacity so that the VM administrator can allocate VVols for block storage. The following example illustrates the task overview.



Use the Add Parity Groups dialog box to add the target parity group to the pool. If necessary, edit capability profiles by modifying the storage container and expanding Thin Image pools or DP pools.

Note: When you expand a Thin Image pool, the configuration information of the storage system changes. To continue tasks using HCS for VVols after the change, you may need to refresh the storage system.

Procedure

1. On the **Resources** tab, select **DP Pools** for the target storage system.
2. Select the target pool, then click **Expand Pools**.
3. Click **Add Parity Groups**.
4. In the **Add Parity Groups** dialog box, select the target parity group, and then click **Add to Pool**.
5. (Optional) Expand a Thin Image pool:
   a. On the **Resources** tab, right-click the target storage system, select **Local Replication**, and then expand the Thin Image pool.
   b. Update the information of the target storage system: On the **Administration** tab, select **Managed Resources**, and then click the **Storage Systems** tab.
6. Click **Show Plan**.
7. Confirm that the details of the plan are correct, and then execute the task.
8. Confirm that the task has completed.
9. Confirm that the capacity of the storage container has increased:
   a. On the **Administration** tab, select **Virtual Disks Management** to display the list of storage containers.
   b. In the **Capacity** area, check that the **Free** space has increased.

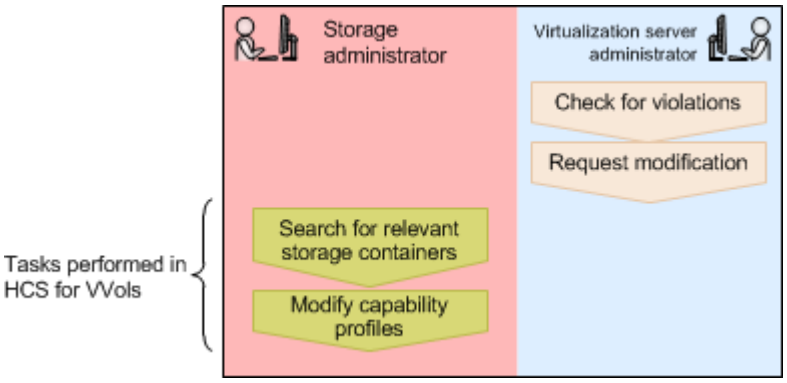
Next steps

If the configuration of the DP pools changes after expanding the capacity, then modify the capability profiles.

Changing a capability profile for block storage VVols

You can change the capability profile information for storage containers when necessary; for example, to correctly allocate VVols after configuration changes have been made, or to correct a capability profile that was incorrectly defined. In the latter case, if the capability profile for the storage container of a virtual machine no longer meets the requirements for the virtual machine, then the VM administrator would search for a storage container that meets the VM storage policy, and would recreate the VVol. Otherwise, if no storage container was found, the VM administrator would ask the storage administrator to modify the capability profile.

The following workflow illustrates the scenarios for modifying a capability profile.



Procedure

1. On the **Administration** tab, select **Virtual Disks Management**, click the **Vendors** tab, and then select **VMware**.
2. On the **VASA Providers** tab, click **VASA Provider (Block Storage)**.
3. On the **Storage Containers and SPBM** tab, select the target storage container from the list, then click **Edit Storage Container**.
4. In the **Capability Profiles** table, select the target capability profile that you want to modify, then click **Define Profile**.
5. Check the applicable information for the **Managed Storage Capabilities** and **Auto-generated Capabilities**.
6. Click **OK** to save the changes and return to the **Edit Storage Container** dialog box.
7. Click **Submit**. The completed operation will be registered as a task.

Next steps

Check the results of the completed task on the Tasks & Alerts tab. Then, after the task has completed, you can check the capability profiles of the modified storage container from the list of Capability Profiles.

Identifying information that exists between VMs and block storage VVols

Within the Administration tab, you can access information about block storage VVols that are associated with the virtual machines in your environment.

There are two main areas of the Administration tab from which you can locate information about block storage VVols:

- When you select a **VASA Provider (Block Storage)** node, click the **VVols** tab.
- When you select a **VASA Provider (Block Storage) Storage Container** node, refer to the table list and table summary for information about LDEV IDs.

In the **VVols** tab, you can see VVol information that is associated with the virtual machine, such as the virtual machine's storage policy which is applied to the VVol, the identifier of the VVol, and the VVol type. In addition, you can click links within the table view to complete the following tasks:

- View details about the selected volume
- Display the **Pool** window
- Display the **Capabilities** dialog, to view the capability profiles and default profiles that are applied to VVols

If at least one VVol is selected, you can click **Edit Labels** to modify the volume label.

When you select a **VASA Provider (Block Storage) Storage Container** node, you can identify the number of available LDEV IDs that can be used to create VVols. You can also view the identifier for the storage container.

Procedure

1. On the **Administration** tab, select the node for which you want to view related VVol information:
   - **VASA Provider (Block Storage)** then click the **VVols** tab.
   - **VASA Provider (Block Storage) Storage Container**
2. The related VVol information that is displayed depends on your selection at step 1.

Deleting storage containers for block storage VVols

When they are no longer needed, you can delete the storage containers that are used by the target virtualization server.

Before you begin

All VVols for block storage that are associated with the storage container must be deleted.

Tip: When all of the DP pools that are included in a storage container are deleted from the resource group, that storage container will be deleted automatically.

Procedure

1. On the **Administration** tab, select **Virtual Disks Management**, click the **Vendors** tab, and then select **VMware**.
2. On the **VASA Providers** tab, click **VASA Provider (Block Storage)**.
3. On the **Storage Containers and SPBM** tab, select the target storage container from the **Storage Containers** list, then click **Delete Storage Container**.
   The **Delete Storage Container** dialog box displays.
4. Click **Submit** to delete the storage container.

Result

The storage container is removed from the Storage Containers list.

Next steps

You can check the results of the completed task on the Tasks & Alerts tab.

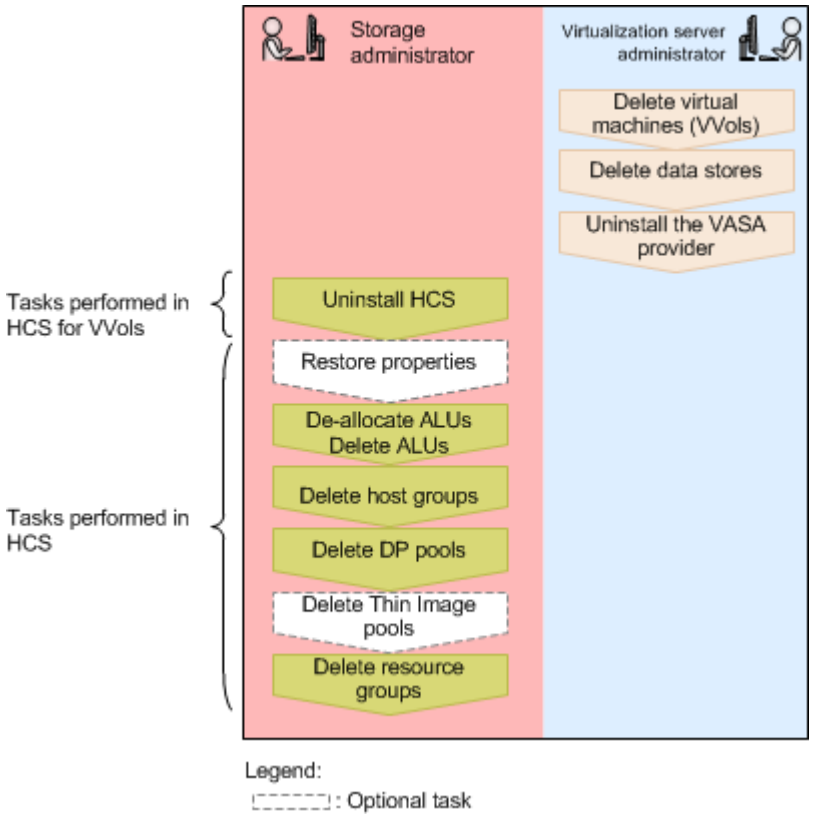
Removing a VVol environment for block storage VVols

When a VMware VVol environment for block storage no longer applies to your site, you can remove it from within Hitachi Command Suite (HCS).

If your site has an obsolete VMware VVol environment, you can remove it by completing the following tasks:

1. **Deleting VVols and datastores, and uninstalling the VASA provider**.
2. **Uninstalling HCS for VVols**. When you uninstall HCS for VVols, the storage containers that are used by the target virtualization server are also removed.
3. **Deleting protocol endpoints**. Remove the allocation of the host group from the protocol endpoints (ALUs) that were used in the target virtualization server. Then, delete the protocol endpoints. You will also remove host groups that are no longer used.
4. **Deleting tenants**. Delete the pools, resource groups, user groups, and users that are no longer used.
5. **Changing HCS system settings**. To continue operations by using HCS (non-VVol version), change the settings so that the operations where VVols are not used are restored.

The following image represents the workflow for removing a VVol environment.



For details about how to return HCS to its initial settings, see the *Hitachi Command Suite Administrator Guide*. For details about how to uninstall HCS, see the *Hitachi Command Suite Installation and Configuration Guide*. For details about how to uninstall a VASA provider, see the manuals for the VASA provider.

Note: The removal of a VMware VVol environment requires that your site coordinates with your VM administrator (for example, to report the deletion of the virtual machines and virtualization server) and with your VASA provider (for example, to report the deletion of the VASA provider).

Deleting a protocol endpoint for block storage VVols

When a VMware VVol environment for block storage no longer applies to your site, you can remove it by deleting its associated protocol endpoints.

You will remove the allocation of the host group from the protocol endpoints that were used in the target virtualization server. Then, delete the protocol endpoints.

You will also remove host groups that are no longer used.

Note: Typically, applicable VVols are deleted before a VMware VVol environment is removed, and before a protocol endpoint is removed because it is no longer required for use in a VMware VVol environment. However, if the

VVols cannot be deleted before performing those tasks (for example, there is some vCenter communication error), then the protocol endpoint will be bound to the VVol. In this case, you must unbind and delete the ALU. For more information, see the *Provisioning Guide* for your storage system.

Before you begin

Delete virtual machines that are associated with the VMware VVol environment.

Procedure

1. On the **Resources** tab, from **Storage Systems**, expand the target storage system tree, right-click **Volumes**, and then select **System GUI**.
2. Specify any required items, then delete the volume.
3. Delete Host Groups:
   a. On the **Resources** tab, click **Storage Systems**, expand the target storage system tree, and then click **Ports/Host Groups/iSCSI Targets**.
   b. In the **Ports/Host Groups/iSCSI Targets** window, click the **Host Groups** tab.
   c. Select the target host group, and then click **Actions**.
   d. Click **Delete Host Groups**.

VVols and file storage

You can use file storage VVols in your environment after completing some initial configuration tasks. Afterward, you can manage your VVol environment for file storage by expanding the storage container capacity, deleting storage containers or removing the VVol environment entirely.

VVol configuration and operation tasks for file storage

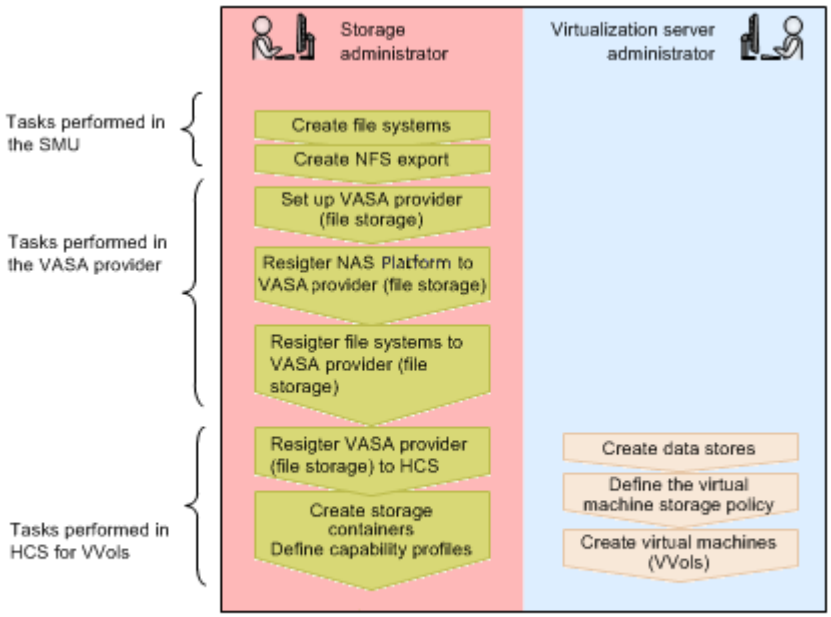
After a separate HCS system has been established for the exclusive use of VVol management (HCS for VVols), you can configure an environment to use VVols for file storage.

Following are the main tasks for configuring your environment to use VVols for file storage:

- Create file systems
- Define an NFS Export
- Register NAS Platform family to the VASA provider (file storage)
- Register the VASA provider (file storage) in HCS for VVols
- Create storage containers

Note: Before you create clones for file storage, the "HNAS file clone on the HNAS" license must be installed.

The following example illustrates the task flow.



For information about setting up HCS for VVols, see the *Hitachi Command Suite Administrator Guide*.

Related tasks

- Changing the VASA provider for file storage VVols on page 428
- Creating storage containers for file storage VVols on page 427
- Modifying storage containers for file storage VVols on page 429
- Deleting storage containers for file storage VVols on page 430
- Expanding the capacity of a storage container for file storage VVols on page 429
- Removing a VVol environment for file storage VVols on page 430

Creating storage containers for file storage VVols

When you create file storage VVols, you will create file systems and define an NFS export on it using the SMU/NAS Manager. Then, you will register a VASA provider for file storage on HCS for VVols (specify the IP address or Host name for the NAS Platform family). Afterward, you can create a storage container and specify a capability profile for it, which ensures that the storage container will match the capabilities of the VM storage policy. When the storage container is created, the VM administrator can use file VVol datastores.

You will use HCS for VVols to create a storage container for file storage VVols.

Procedure

1. Launch HCS for VVols.
2. From the **Administration** tab, click **Virtual Disk Management**, then select **VMware**.
3. Confirm that no VASA provider for managing file VVols has been registered, then click the **Add VASA Provider** button.
4. Enter the IP address or Host name for the VASA provider (file storage).
5. Confirm that the VASA provider (file storage) is registered in the list of VASA providers, then select the registered **VASA Provider (File Storage)** and open the **Administration** window.
6. Click **Create Storage Container**.
7. Specify basic information for the storage container:
   a. In the **Name** field, enter the name of the storage container.
   b. (Optional) In the **Description** field, enter a description.
   c. From the **File System** list, select the target storage system.
8. Define a capability profile for file systems that are listed in the table for which the **Name** column indicates **Undefined**:
   a. Click **Define Profile** to create a new profile.
   b. In the **Name** field, enter the name of the capability profile.
   c. (Optional) In the **Description** field, enter a description.
   d. Within the **Managed Storage Capabilities** sections, either specify settings or check existing settings to ensure that the storage container to be created will match the capabilities of the VM storage policy.
   e. Click **OK** to save the new capability profile and return to the **Create Storage Container** dialog box.
9. Click **Submit**.
   The completed operation will be registered as a task.

Changing the VASA provider for file storage VVols

The VASA provider for file storage VVols must be registered in HCS for VVols. After the registration has been done, you can change the information that is associated with the VASA provider.

When necessary, you can change information about a registered VASA provider, such as the associated IP address (or, host name).

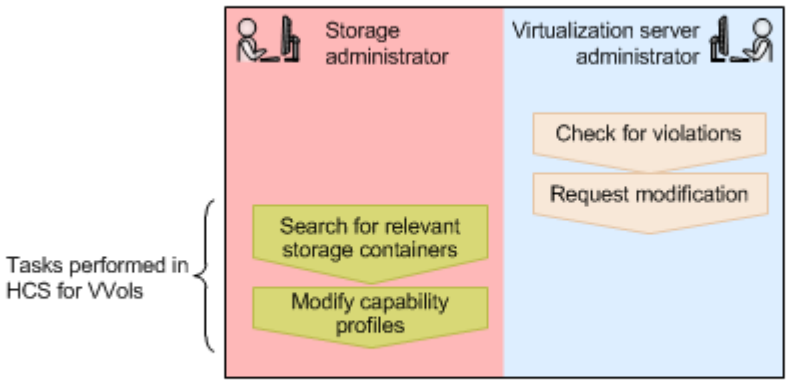
Procedure

1. On the **Administration** tab, navigate to **Virtual Disks Management**, click the **Vendors** tab, and select **VMware**.
2. On the **VASA Providers** tab, click **Edit VASA Provider**.
3. Modify the applicable information.
4. Click **OK**.

Modifying storage containers for file storage VVols

You will use HCS for VVols to change information that is associated with an existing storage container.

The following workflow illustrates the scenarios for initiating a change to a storage container.



Procedure

1. On the **Administration** tab, navigate to **Virtual Disks Management**, click the **Vendors** tab and then select **VMware**.
2. Click the **VASA Providers** tab, select **VASA Provider (File Storage)**, and then click the **Storage Container and SPBM** tab.
3. Click **Edit Storage Container**.
4. Modify the applicable information:
   a. In the **Name** field, enter a different name for the storage container.
   b. (Optional) In the **Description** field, enter a description.
   c. From the **File System** list, select a different target storage system.
5. Click **Submit**.
   The completed operation will be registered as a task.

Expanding the capacity of a storage container for file storage VVols

When a file storage system exists, you can modify a storage container to increase its capacity.

Based on the workflow that has been established for your site, you will either receive an e-mail notification when a threshold is exceeded for a file system, or you will be notified by the VM administrator that a capacity shortage exists. Depending on the case, you can either expand the applicable file

system by modifying the associated storage container, or you can increase the free space by adding a new file system to the storage container.

If you need to create additional file systems, then you can add them using the System Management Unit (SMU)/NAS Manager. Afterward, you will modify the storage container using HCS for VVols.

Procedure

1. Launch the SMU/NAS Manager.
2. Based on the type of capacity shortage notification that your site has received, complete one of the following tasks:
   - Expand the affected file system by modifying the associated storage container.
   - Increase free space by adding a new file system to the storage container. Proceed to step 3.
3. Launch HCS for VVols.
4. Click the **Administration** tab, select **Virtual Disk Management**, then click **VMware**.
5. Select **VASA provider (File Storage)** and locate the obsolete storage container.
6. Click **Edit Storage Container**.
7. Modify the capability profile.

Deleting storage containers for file storage VVols

If your site no longer has a use for an existing storage container, you can remove it.

You will use HCS for VVols to locate the obsolete storage container, and delete it.

Procedure

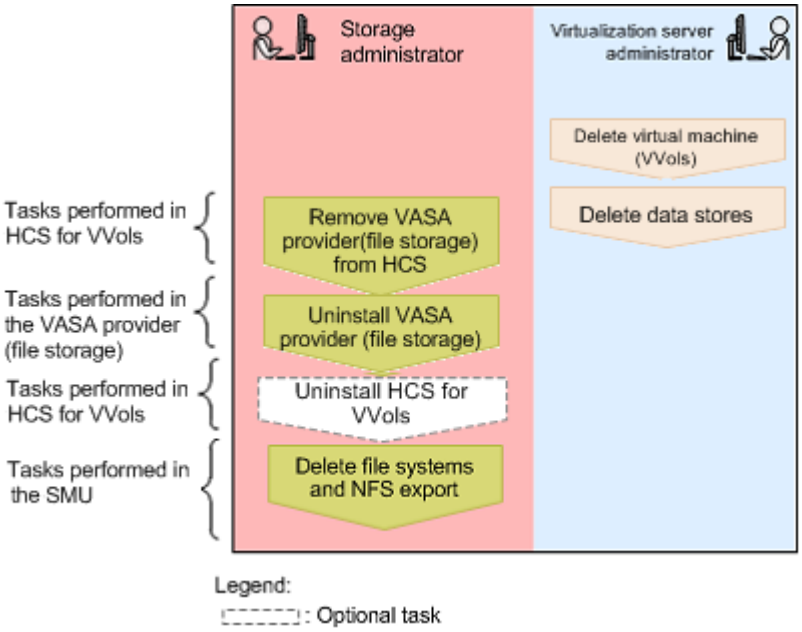
1. Launch HCS for VVols.
2. From the **Administration** tab, click **Virtual Disk Management**, then select **VMware**.
3. Click the **VASA providers** tab and select **VASA Provider (File Storage)**.
4. From the **Storage Container and SPBM** tab, locate the obsolete storage container, then click **Delete Storage Container**.

Removing a VVol environment for file storage VVols

You will use HCS for VVols and the VASA provider to remove a file storage VVol environment.

If your site no longer requires a file storage VVol environment, you can choose to remove it entirely, or you can continue to retain the VVol environment for block storage.

The following workflow illustrates the required and optional steps that are involved in removing a file storage VVol environment.



Procedure

1. In HCS for VVols, consider your site's current VVol environment needs, then complete one of the following options:
   - If your site will no longer use a VVol environment for block storage, uninstall HCS for VVols.
   - If your site will continue to use a VVol environment for block storage, remove the VASA provider for file storage from HCS for VVols.
2. Launch the SMU/NAS Manager.
3. Delete file systems.
4. Delete the specified **NFS Export** information.

11

Optimizing storage performance

This module describes how to improve your storage.

- About optimizing storage
- About optimizing HBA configurations
- About high temperature mode
- Managing cache logical partitions
- Data mobility
- Data migration

About optimizing storage

Hitachi Command Suite allows you to manage storage by allocating volumes, expanding tiers and DP pools, and performing migration, based on information acquired from checking summaries, alerts, performance statistics, and the operating status of storage resources. Storage utilization can also be improved with effective management of HBA and cache (CLPRs) resources.

Using HDT pools, you can manage performance and capacity to optimize storage by creating parity groups in the following dynamic tiers:

- Tier 1 using parity groups for best performance
- Tier 2 using parity groups for next best performance
- Tier 3 using parity groups for capacity independent of drive type or RAID level

Common storage optimization issues include:

- If the used capacity for a DP pool has reached or exceeded its threshold, or a volume that satisfies a capacity requirement cannot be created or assigned because of insufficient unused capacity, add DP pool volumes to increase the capacity of the DP pools. If the capacity of a specific drive is insufficient when using an HDT pool, increase the capacity by mixing different drive types or RAID levels in Tier 3.
- If the usage rate of the file system has reached the threshold value, expand the file system to increase the capacity that can be allocated.
- If C/T delta values have degenerated and reached the threshold, use the Replication tab to confirm the degeneration factor and countermeasures, and use Device Manager, Replication Manager or Storage Navigator to resolve the problem.
- If the performance of a DP pool has decreased and data I/O is slow, add more DP pool volumes to distribute loads within the DP pools. Another option is to perform volume migration to distribute I/O loads on the DP pools.
- When using the HDT pool, performance problems may occur at certain times. Ensure that monitoring occurs during periods when I/O loads are occurring. You can:
  - Start or stop the monitoring/relocation process manually in accordance with known times for load changes.
  - Cancel the monitoring process during periods of low activity.
- If HDT volume applications switch between online and batch processing, it can be helpful to save optimized volume data placements, by processing method, as profiles. By applying the corresponding profile before beginning processing, the data placement that fits the characteristics of the processing method is restored.
- When using HDT pools, you want to prioritize the data relocation of HDT volumes for which capacity and access patterns vary widely, but I/O

operations decrease without relocating effectively. You can disable the relocation of HDT volumes for which the current data location presents no problems, to reduce relocation load.

- When using HDT pools, important data is allocated to the lower hardware tier because it has fewer I/O accesses. To prevent unwanted relocations, set a specific hardware tier for the HDT pool by configuring tiering (Tier 1, Tier 2, and Tier 3).
- When using HDT pools, use the flexibility of tiering to spread the data in a host volume across multiple layers of parity groups (high-speed, next highest speed, and low-speed) that are contained in a pool structured for this purpose.
- When using HDT pools, understand that if different drive types and/or RAID levels are mixed in a single tier, they will all be considered equal for data placement regardless of page access frequency. As a result I/O performance will be dependent on the drive type characteristics and RAID level on which any given page resides.
- If the load on a volume is too high when volume data is backed up to tape, create a copy pair for the volume. Then do a tape backup by using the copied volume (as a secondary volume).
- If it is not possible to assign a high-performance volume to a host because all unassigned volumes are low performance, perform volume migration so that less frequently used data is migrated to a low-performance volume, and then assign the now available high-performance volume to an appropriate host.
- If the usage frequency of an application increases, you can add an HBA and increase LUN paths to improve data transmission performance and throughput requirements.

Related concepts

- About data mobility on page 444
- About data migration on page 451
- About optimizing HBA configurations on page 435

Related tasks

- Creating a CLPR on page 439

About optimizing HBA configurations

Using HCS, optimize or maintain server HBA configurations in support of high availability and performance/throughput requirements.

The initial allocation of volumes to a server typically occurs with the allocate volumes dialog. LUN paths are established at this time. Host groups are used to control access to ports and volumes, meaning all the hosts in the group are using the same ports to access allocated volumes for the hosts.

Over time, heavily used servers might exhibit the need for improved high-availability, and/or improved I/O and throughput performance. HCS provides for the optimization and maintenance of server HBA configurations, as follows:

- For optimizing a server HBA configuration, you can add one or more HBAs to a server, and simultaneously identify one or more HBAs (WWNs) in the host group for the purpose of inheriting existing LUN path information for the newly added HBAs. This provides a fast and easy way to add HBAs, for example increasing from one HBA to two HBAs, or two HBAs to four HBAs.
- In terms of maintaining current performance levels, you can remove a failed HBA and add a new HBA, or you can add a new HBA then remove the failed HBA after the new HBA is in service.

Note that redundant HBAs can provide improved high availability, performance and throughput for a server, unless the server itself fails. The solution for server failure is clustering.

Related tasks

- Adding an HBA on page 436

Adding an HBA

Add an HBA to improve performance and throughput requirements. When adding an HBA, specify the WWN of the new HBA and then select a WWN of an existing HBA from which to model paths.

Before you begin

- Identify the new WWN for the HBA that is being added
- Identify the WWN from which to model paths
- Verify that the new HBA is physically connected

Procedure

1. On the **Resources** tab, select **Hosts**.
2. After selecting the target operating system, select the target host row and click **More Actions > Add HBAs**.
3. Enter the **New WWN** or select a WWN from the list.
4. Enter the **WWN from which to model paths** or select a WWN from the list.
5. Click **Add.**
6. In the WWN Pairs list, verify that the listed HBA WWN combination are paired correctly.



Tip:

- If the WWN information is updated when the host is refreshed, the target WWN might not be displayed in the list. In this case,

you need to manually enter the WWN of the HBA you are adding.
- To edit a WWN nickname from the list of **WWN Pairs**, click **Edit WWN Nicknames**.

7. Click **Show Plan** and confirm that the information in the plan summary is correct. If changes are required, click **Back**.
8. (Optional) Update the task name and provide a description.
9. (Optional) Expand **Schedule** to specify the task schedule.
   You can schedule the task to run immediately or later. The default setting is **Now**. If the task is scheduled to run immediately, you can select **View task status** to monitor the task after it is submitted.
10. Click **Submit.**
    If the task is scheduled to run immediately, the process begins.
11. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.

Result

When the task completes, the new WWN is added and LUN path settings are established to the host.

Related tasks

- Editing LUN paths when exchanging a failed HBA on page 243
- Editing LUN paths when adding or exchanging an HBA on page 244

About high temperature mode

You can use Hitachi Command Suite to enable high temperature mode, which is a licensed feature that allows the VSP G1000, VSP G1500, and VSP F1500 storage systems to run at higher temperatures (60.8°F to 104°F / 16°C to 40°C), saving energy and cooling costs.

When high temperature mode is disabled, the VSP G1000, VSP G1500, and VSP F1500 storage systems run at standard temperatures (60.8°F to 89.6°F / 16°C to 32°C).

Normal and high temperature alerts

When high temperature mode is disabled, an alert displays when the temperature in the storage system exceeds 89.6°F / 32°C.

When high temperature mode is enabled, an alert displays when the temperature in the storage system exceeds 104°F / 40°C.

Caution: Before you enable high temperature mode, review the *Hardware Guide for Hitachi Virtual Storage Platform G1000, G1500, and F1500* for restrictions and important information.

Related references

- Enabling high temperature mode for VSP G1000, VSP G1500, and VSP F1500 storage systems on page 438

Enabling high temperature mode for VSP G1000, VSP G1500, and VSP F1500 storage systems

For VSP G1000, VSP G1500, and VSP F1500 storage systems, you can enable high temperature mode.

Before you begin

- You must install a valid license for this feature.

Caution: Before you enable high temperature mode, see the *Hardware Guide for Hitachi Virtual Storage Platform G1000, G1500, and F1500* for important information.

Procedure

1. On the **Resources** tab, click **Storage Systems**, and then expand **All Storage Systems** and the target storage system.
2. Click **Components**.
3. Click **Edit High Temperature Mode**
4. Click **Enable (16-40 degrees C)**.
5. Click **Finish.**
6. In the **Confirm** window, verify the settings and enter a task name.
7. Click **Apply** to register the task. If the **Go to tasks window for status** check box is checked, the **Task** window opens.

Result

After the task completes, high temperature mode is enabled.

Related concepts

- About high temperature mode on page 437

Managing cache logical partitions

This module describes how to manage cache logical partitions (CLPR), including managing the assignment of resources to the CLPR.

If you are using a VSP G400, VSP G600, or VSP G800 storage system with a NAS module and the firmware version is earlier than 83-03-25-XX/XX, do not use the CLPR named "NASSystemCLPR" for operations such as volume allocation, virtualization, or data migration. If you use this CLPR for such operations, system performance might be severely affected.

Creating a CLPR

You can create partitioned cache as a means of providing predictable performance levels for server applications, and providing memory protection.

Caution: Creating CLPRs can significantly degrade host performance and should be performed during the initial installation and setup of the storage system or during a maintenance window. Before creating a CLPR, read *Cautions and restrictions for Virtual Partition Manager*.

If no CLPRs have been created, the entire cache is displayed as CLPR0. When the first CLPR is created, CLPR1 is added. Up to CLPR31 can be created.

The default cache capacity is 8 GB. CLPRs can be created by assigning the necessary capacity from CLPR0. If you use Cache Residency, the remaining cache capacity after subtracting the Cache Residency capacity from the cache capacity of CLPR0 must be at least 8 GB.

Procedure

1. On the **Resources** tab, expand the **Storage Systems** tree, and select the target storage system.
2. Choose one of the following options.
   - For VSP Gx00 models and VSP G1000, VSP G1500, and VSP F1500 storage systems:
     Select **Cache Partitions**.
   - For other available storage systems:
     From the **Actions** list in the application pane, select **Element Manager**. Refer to the documentation for the native management tool for your storage system.
3. On the **Cache Partitions** tab, click **Create CLPRs** to open the **Create CLPRs** window. **CLPR ID** displays the first available CLPR number or a blank if no CLPR number is available.
4. In **CLPR Name**, enter a CLPR name (maximum 16 alphanumeric characters). Each CLPR name must be unique. You cannot use a CLPR name that is already reserved.
5. In **Total Cache Size** select the cache capacity.
   The default size is 8 GB. You can select 8 GB or a higher value in increments of 4 GB. The maximum value is 504 GB (by subtracting 8 GB from the cache capacity of the storage system), but the maximum available capacity (by subtracting the total usage capacity of other CLPRs from the total capacity of the storage system) is displayed as the upper limit value.

6. In **Resident Cache Size** select the resident cache capacity.

   The default is 0 GB, and you can select 0 GB or higher value in increments of 0.5 GB. The maximum value is 504 GB (cache residency capacity of the storage system), but the maximum available capacity (by subtracting the total usage capacity of other CLPRs from the total capacity of the storage system) is displayed as the upper limit value.
7. In **Number of Resident Extents**, enter the number of resident cache.

   The default is 0, and you can specify 0 to 16384. The maximum available capacity (by subtracting the total usage capacity of other CLPRs from the total capacity of the storage system) is displayed as the upper limit value.
8. Click **Add**. The created CLPR is added to the **Selected CLPRs** table.

   To delete a CLPR from the **Selected CLPRs** table, select the CLPR and click **Remove**. To change the settings of an existing CLPR, select the CLPR and **Change Settings** to open the **Change Settings** window.
9. Click **Finish**.
10. Check the settings in the **Confirm** window, enter the task name in **Task Name**, and click **Apply**.

Result

The CLPR is created. A newly created CLPR has no resources (parity groups). To migrate resources to the new CLPR, see *Migrating resources to and from a CLPR*.

Related tasks

- Migrating resources to and from a CLPR on page 440
- Editing the settings of an existing CLPR on page 441
- Adjusting the cache capacity of a CLPR on page 443
- Deleting a CLPR on page 443

Migrating resources to and from a CLPR

After creating a CLPR, you can migrate resources (parity groups) from existing CLPRs to the new CLPR. Before deleting a CLPR, you must first migrate resources that you want to keep to other CLPRs.



Caution: Migrating resources to and from CLPRs can significantly degrade host performance and should be performed during the initial installation and setup of the storage system or during a maintenance window.

When migrating resources to and from CLPRs:

- You can migrate resources only within the same CU
- All interleaved parity groups must be in the same CLPR
- If a parity group contains one or more LDEVs that have defined Cache Residency Manager extents, you cannot migrate that parity group to another CLPR.

Procedure

1. On the **Resources** tab, expand the **Storage Systems** tree, and select the target storage system.
2. Choose one of the following options.
   - For VSP Gx00 models and VSP G1000, VSP G1500, or VSP F1500 storage systems:
     Select **Cache Partitions**.
   - For other available storage systems:
     From the **Actions** list in the application pane, select **Element Manager**. Refer to the documentation for the native management tool for your storage system.
3. In **CLPR** on the **Migrate CLPR Resources** window, select the CLPR number to which the resource is assigned.
4. In **CU**, select the CU to which the resource is assigned.
5. In **Available Parity Groups** or **Available Virtual Volumes**, select the resource that you want to migrate, and then click **Set>**. This adds the selected resource to the **CLPR** table.
   To check the information of the CLPR, select the CLPR and click **Detail** to open the **CLPR Properties** window.
6. When you are done adding resources to migrate, click **Finish**.
7. On the **Confirm** window, check the settings carefully.
8. Click **Finish**.
9. When you are ready to start the resource migration, enter the task name in the **Confirm** window, and then click **Apply**.

Result

The specified resources are migrated.

Related tasks

- Editing the settings of an existing CLPR on page 441
- Creating a CLPR on page 439

Editing the settings of an existing CLPR

You can change the following settings after creating a CLPR:

- CLPR name
- Total cache size
- Resident cache size
- Number of resident cache extents

Procedure

1. On the **Resources** tab, expand the **Storage Systems** tree, and select the target storage system.
2. Choose one of the following options.

- For VSP Gx00 models and VSP G1000, VSP G1500, and VSP F1500 storage systems:
  Select **Cache Partitions**.
- For other available storage systems:
  From the **Actions** list in the application pane, select **Element Manager**. Refer to the documentation for the native management tool for your storage system.

3. On the **Cache Partitions** tab, select the CLPR to be edited, and then click **Edit CLPR** to open the **Edit CLPR** window.
4. Edit the CLPR settings as follows. The CLPR number cannot be changed.
   a. In **CLPR Name**, edit the CLPR name as needed. You can use up to 16 alphanumeric characters but you cannot specify CLPR names that are already reserved.
   b. In **Total Cache Size**, edit the CLPR cache capacity as needed. The minimum cache size is 8 GB, and the size can be changed in increments of 4 GB. The maximum available capacity (by subtracting the total usage capacity of other CLPRs from the total capacity of the storage system) is displayed as the maximum size. The remaining cache capacity (by subtracting the resident cache size from the total cache size) must be 8 GB or more. Ensure that the remaining capacity in CLPR0 is at least 8 GB when selecting the total cache size or resident cache size.
   c. In **Resident Cache Size**, edit the resident cache capacity as needed. The minimum size is 0 GB, and the size can be changed in increments of 0.5 GB. The maximum available capacity (by subtracting the total usage capacity of other CLPRs from the total capacity of the storage system) is displayed as the maximum size. If resident cache exists in the selected CLPR and you select a resident cache size that is smaller than the existing resident cache, an error will occur. Before decreasing the size, check the resident cache capacity set in the selected CLPR.
   d. In **Number of Resident Extents**, edit the number of resident cache extents as needed. You can specify from 0 to 16384. The maximum available capacity (by subtracting the total usage capacity of other CLPRs from the total capacity of the storage system) is displayed as the maximum value. If resident cache exists in the selected CLPR and you select fewer resident cache extents than the existing number of extents, an error will occur. Before decreasing the number, check the number of resident extents set in the selected CLPR.
5. When you are done editing the CLPR settings, click **Finish**.
6. On the **Confirm** window, check the settings carefully.
7. When you are ready to change the settings, enter the task name in the **Confirm** window, and then click **Apply**.

Result

The specified resources are changed.

Related tasks

- Adjusting the cache capacity of a CLPR on page 443

- Creating a CLPR on page 439

Adjusting the cache capacity of a CLPR

You need to adjust the cache capacity of a CLPR when Cache Residency Manager operations decrease the cache capacity of the CLPR.

Procedure

1. Cancel the Cache Residency Manager bind mode or priority mode setting.
2. Edit the settings of an existing CLPR to change the cache capacity of the CLPR.
3. Set the bind mode or priority mode again.

Result

The cache capacity is adjusted.

Related tasks

- Creating a CLPR on page 439
- Editing the settings of an existing CLPR on page 441

Deleting a CLPR

You can only delete CLPRs that you created. CLPR0 cannot be deleted.

Caution: Deleting a CLPR deletes all resources and data in the CLPR. Before deleting a CLPR, migrate all resources that you want to keep (for example, parity groups) to another CLPR that will not be deleted.

Procedure

1. On the **Resources** tab, expand the **Storage Systems** tree, and select the target storage system.
2. Choose one of the following options.
   - For VSP Gx00 models and VSP G1000, VSP G1500, and VSP F1500 storage systems:
     Select **Cache Partitions**.
   - For other available storage systems:
     From the **Actions** list in the application pane, select **Element Manager**. Refer to the documentation for the native management tool for your storage system.
3. On the **Cache Partitions** tab, select the CLPR that you want to delete, and then click **Delete CLPRs** to open the **Delete CLPRs** window.
4. In the **Selected CLPRs** table, check the CLPR that you want to delete. If you want to cancel the deletion, click **Cancel** to go back to the **Cache Partitions** tab.
5. Click **Apply**.

Result

The CLPR is deleted.

Related tasks

- Migrating resources to and from a CLPR on page 440

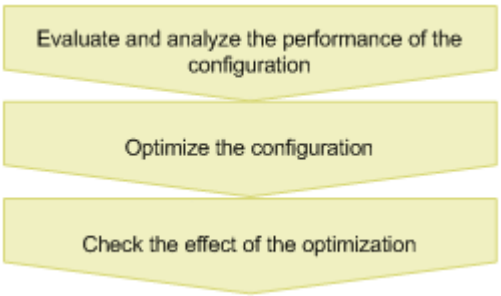
Data mobility

An administrator can obtain information required to see the operating status of storage resources and optimize data placement.

About data mobility

The Mobility tab provides information required to see the operating status of storage resources and optimize data placement.

To efficiently use a storage resource, use the Mobility tab to optimize the data arrangement. In the Mobility tab, information used to determine whether the data arrangement is appropriate is collectively displayed. In addition, you can use the Mobility tab to centrally perform operations regarding moving data such as data migration. To use the Mobility tab, you must first register a Tiered Storage Manager license. Using Mobility tab information, you can evaluate and analyze whether the hardware tiers are being used appropriately and whether the volumes are providing the expected performance. Based on these results, you can optimize data placement by migrating data and setting up an HDT tiering policy (which specifies the hardware tiers in which volumes are to be placed). After completing these operations, check their effect from the Mobility tab and then perform any other necessary data optimization operations. By repeating this process, the placement of data can be optimized. The following figure shows the process of optimizing data placement.



From the Mobility tab, data placement can be optimized from either the application (logical group) perspective or from the DP pool perspective. To respond to requests from host or application administrators, or to confirm whether HDT pools and storage resources are being used according to the needs of each application, check the balance between the performance and the configuration ratio of logical groups. To regularly monitor storage resources, or if there is a problem in the analysis of logical groups, check the

balance between the performance and the configuration ratio of each DP pool. If there is a problem, you must check the operating status of each volume in the DP pool and then optimize the configuration.

When periodically monitoring logical groups or DP pools during routine work, check the ranking reports in the Mobility tab so you can efficiently find logical groups or DP pools that have problems in their operational status.

The Mobility tab can be used to check the total capacity and used capacity rate for each logical group, DP pool, or each volume in a logical group or DP pool. As with HDT pools, you can also check which volumes are using the largest proportions of each hardware tier. In addition, the following performance information can also be viewed:

- Quantity of I/O operations compared to the calculated maximum number of I/O operations

In addition, if linked with Hitachi Tuning Manager, the following performance information can also be viewed:

- Throughput (IOPS) for DP pools, volumes, and hardware tiers
- Past volume performance data (to look for trends)
- Response time of the volume
- Busy rate for parity groups

Note: It's not possible to determine the load status from the busy rate for parity groups for flash drives (FMD, FMC, and SSD) of storage systems except for VSP G1000 (microcode version is 80-05-01-XX/XX or later), VSP G1500, and VSP F1500. Because of the features of the flash drives, larger values might be displayed regardless of the load status.

For details on linking to Hitachi Tuning Manager, see the *Hitachi Command Suite Administrator Guide*.

Note: If Hitachi Tuning Manager is not linked, or information cannot be obtained due to storage system type, or the Tuning Manager - Agent for RAID software version, hyphens are displayed in performance information columns.

From the Mobility tab, the following operations can be performed to optimize data placement:

- Configuring and editing of an HDT tiering policy
- Migrating volumes
- Manually starting or stopping the monitoring and relocation of HDT pools
- Editing tier relocation of HDT volumes
- Expanding, shrinking, or changing settings for DP pools
- Expanding DP volumes and reclaiming zero pages
- Referencing trend charts, which indicate capacity data and performance data for DP pools, and exporting data to CSV files

By creating tiers of volumes or DP pools based on use or performance, you can migrate data to volumes or DP pools that match the relevant conditions simply by specifying a tier as the migration target.

Related tasks

- Optimizing data placement in a logical group on page 449
- Optimizing data placement in a DP pool on page 449

Reports for logical groups displayed in the Mobility tab

Several reports are available on a logical group basis for monitoring performance related items such as response time and unused capacity.

You can display one of the following ranking reports for a logical group in the Mobility tab by selecting a top folder of the Public or Private folder, and a top-level folder for each user group under the Private Logical Group folder:

- **Bottom 10: Logical Groups with High Response Time**
  Checks for logical groups including volumes whose response time is slow. Used to determine logical groups which do not satisfy required performances. This ranking is displayed only when you have the license for Hitachi Tuning Manager.
- **Top 10: Logical Groups with Lowest Unused Capacity (Excluding S-VOLs of Casading Thin Image Pairs)**
  Checks for logical groups that have little unused capacity left for the allocated DP volumes. Used to determine logical groups which do not have enough capacity.
  Because the space in Thin Image S-VOLs that support cascade configurations will never be used, the space is not included in the unused capacity.
- **Top 10: Logical Group Used Capacity by Drive Types**
  Checks for logical groups that use a large amount of capacity by drive type. Used to determine whether an unexpected volume is occupying a high-speed drive or logical groups which have invalid data arrangements.

The following summary information is displayed when each logical group is selected:

- **Summary**
  Provides summarized information in a list, such as the number of volumes and the total capacity of a logical group, or the last update date. Based on this information, you can review the usage status of the storage resources in the logical group.
- **Drive Type Summary**
  Checks whether the logical group has the necessary capacity. Checks the total capacity of the volumes and the used capacity of each drive type within the logical group.
- **Drive Type Distribution**

Checks how many of each drive type is used for each DP pool that is included in a logical group. Used to check whether the configuration of the hardware tier is appropriate.

Related concepts

- About data mobility on page 444

Reports for DP pools displayed in the Mobility tab

The Mobility tab displays several types of reports for DP pools related to performance and capacity utilization.

Ranking reports for DP groups displayed in the Mobility tab are as follows:

- **Top 10: Pools with the Highest Used %**
  Check on a regular basis for DP pools with a high used capacity rate to determine DP pools that need pool expansion or data migration.
- **Top 10: Pools with the Highest Subscription (Excluding S-VOLs of Cascading Thin Image Pairs)**
  Check on a regular basis for DP pools with a high reservation rate to determine DP pools that need pool expansion or data migration on a regular basis.
  Because the space in Thin Image S-VOLs that support cascade configurations will never be used, these volumes are not counted as reserved volumes.
- **Top 10: Tiers with the Highest Performance Utilization**
  Check on a regular basis for hardware tier configurations with high ratio of actual I/O operations to the calculated maximum number of I/O operations. Used to determine HDT pools that include hardware tier configurations with saturated performance.

The information displayed when each DP pool is selected is as follows:

- **Summary**
  Summarizes information in a list, such as the total capacity and usage of DP pools, and the reservation rate. Based on this information, you can review the usage status of the DP pool.
  In addition, you can determine whether you need to add capacity to the DP pool by selecting Show Trend Chart from the Actions menu, and then checking each pool's trends in the increase of IOPS and in capacity usage. You can change the display period of the trend chart as needed, or save the trend chart as a report after exporting it to a CSV file.

If an HDT pool is selected, the following information can also be viewed:

- **Tier Summary**
  Check information such as drive type, drive rotational speed, and capacity usage rate for each hardware tier. Based on this information, you can review the usage status of hardware tiers. Note that this summary only appears for HDT pools.
- **Capacity & Performance Utilization**

Check the total capacity and used capacity of the HDT pools for each hardware tier. This lets you confirm whether the HDT pool has the necessary capacity and identifies which hardware tier needs more capacity added to it. This summary only appears for HDT pools. For example, if capacity is increasing, but IOPS shows little change, then adding more capacity for future use is suggested with this capacity added to Tier 2 or Tier 3. If capacity shows little change, but IOPS is increasing, then adding more capacity for future use is suggested with this capacity added to the upper tiers (Tier 1 and Tier 2), or reducing Tier 3 by the amount needed. If both capacity and IOPS are increasing, then adding more capacity for future use is necessary. You can determine to which tiers this capacity is added based on the trend graph and distribution chart for each tier.

- **Show Trend Chart**
  Check the trend chart, which combines and displays the capacity usage and performance information for each hardware tier. Based on this information, you can determine whether the hardware tier configuration in the HDT pool is appropriate, and which hardware tiers need additional capacity. You can change the display period of the trend chart as needed, or save the trend chart as a report after exporting it to a CSV file. If the display period of the trend chart is seven days or more, the chart displays the average of the obtained information for each day.
- **Capacity Utilization**
  Check the total capacity and used capacity of the HDT pools for each hardware tier. Based on this information, you can check whether the HDT pool has the necessary capacity. You need to perform capacity planning for pools in which the remaining capacity is small. Look at the trends for how pool contents increase, and estimate when you must increase capacity, and possibly add additional media.
- **Performance Utilization**
  Check the ratio of actual I/O operations to the calculated maximum number of I/O operations for each hardware tier. Based on this information, you can determine the hardware tiers whose performance is in a saturated state that then require an optimization operation (such as editing the tiering policy or data migration).
- **Top 10: Volumes Sharing Tier Capacity**
  Check the usage ranking of volumes that have the largest use capacity in each hardware tier. Based on this information, you can check whether data placement is appropriate, for example by checking whether an unexpected volume is occupying a high-speed drive. This summary only appears for HDT pools.
- **View Tier Properties**
  Check the page access distribution per average number of I/O operations, which is based on the monitoring information of HDT pools.
  Based on this information, you can check whether the hardware tier configuration or tiering policy settings are appropriate.

From the page access distribution for HDT pools, you can narrow the range of displayed items to only tiering policies or only HDT volumes, as required.

Related concepts

- About data mobility on page 444

Optimizing data placement in a logical group

Data placement can be optimized in either of the supported logical group types (Public Logical Groups or Private Logical Groups).

Before you begin

- A registered Tiered Storage Manager license.
- A registered Hitachi Tuning Manager license (when referencing performance information). If Hitachi Command Suite is linked with Hitachi Tuning Manager, you can display the Historical Report and check performance information by clicking Analyze Performance. For information about changing the report window settings for the Historical Report, see the *Hitachi Command Suite Tuning Manager User Guide*.

Procedure

1. From the tree view of the **Mobility** tab, select **Logical Groups**.
2. From **Logical Groups**, select either **Public Logical Groups** or **Private Logical Groups**, and in the logical group ranking, check for logical groups that are responding late, have a small unused capacity, or have a large used capacity for each drive type (for example: FMD, FMC, SSD, SAS, or SATA).
3. If the capacity or performance of the entire logical group is not satisfactory, click the link of the logical group name to check the total capacity of the volumes and the used capacity of each drive type (for example: FMD, FMC, SSD, SAS, or SATA) within the logical group.
4. If the balance between the logical group performance and configuration ratio is not satisfactory, select the row of the target volume, and click **Migrate Data**, **Edit Tiering Policy**, or one of the other buttons to remedy the problem.
5. Select **Logical Groups** from the **Mobility** tab to check optimization of the data placement.

Related concepts

- About data mobility on page 444

Related tasks

- Optimizing data placement in a DP pool on page 449

Optimizing data placement in a DP pool

Data placement can be optimized in a DP pool.

Before you begin

- A registered Tiered Storage Manager license.
- A registered Hitachi Tuning Manager license (when referencing performance information). If Hitachi Command Suite is linked with Hitachi Tuning Manager, you can display the Historical Report and check performance information by clicking Analyze Performance. For information about changing the report window settings for the Historical Report, see the *Hitachi Command Suite Tuning Manager User Guide*.
- Ensure that HDT pools are being monitored (when referencing Performance Utilization, IOPS for a hardware tier, or View Tier Properties).

Procedure

1. From the tree view of the **Mobility** tab, select **DP Pools**.
2. Select **DP Pool** to check for DP pools whose used capacity rate, reservation rate, or performance usage rate is high.
3. If a DP pool does not meet the expectations for capacity or performance, click the link of the DP pool in the ranking.
4. Check the usage status and performance of each hardware tier in the DP pools by performing the following:
   - To check if the usage capacity of the HDP pool is appropriate, from the **Actions** menu, select **Show Trend Chart**, and then check the IOPS and usage capacity of each pool.
   - To check if the configuration and performance of each hardware tier in the HDT pool is appropriate, from the **Actions** menu (or from the **Capacity & Performance Utilization** link), select **Show Trend Chart**, and then check the capacity and performance of each hardware tier. To save the contents of the Trend Chart as a report, click **Export to CSV** and save the contents as a CSV file to a location of your choice.
   - To check page access distribution per average number of I/O operations, from the **Actions** menu, select **View Tier Properties**.
5. Select the DP volumes that have a problem and optimize the data placement by clicking:
   - **Migrate Data**
   - **Edit Tiering Policy**
   - **Expand DP Volume**, or one of the other buttons
6. On the **Mobility** tab, select **DP Pools** to confirm data placement optimization.

Related concepts

- About data mobility on page 444

Related tasks

- Optimizing data placement in a logical group on page 449

Data migration

This module describes data migration.

About data migration

Data migration involves moving data from a set of source volumes to a set of destination volumes or DP pools, including those in externally-connected storage systems.

For example, you can migrate data when your volume performance decreases or your environment changes. When relocating data based on the results of performance monitoring for volumes being used by servers and applications, or when migrating data from an old storage system to a newly set-up storage system, the data in multiple volumes can be migrated to other volumes or DP pools.

The source volume can be selected by host, file server or by logical group. The source volume also can be selected from some of the volumes allocated to a host, file server, some of the volumes displayed in the list of volumes in the file server, some of the volumes in a logical group, some of the volumes in a tier. Migration source volumes can also be selected from volume search results. If a volume that cannot be migrated is included in the selected volumes, that volume is excluded from migration and only volumes that can be migrated, will be migrated.

Migration targets are selected from tiers, DP pools, or resource search results. If DP pools are selected, DP volumes with the same capacity as the migration-source volumes are automatically created. The existing DP volumes also can be specified as migration targets.

Note the following:

- When the migration source and the migration target are selected, individual pairs of migration source volumes and migration target volumes (or DP pools) are automatically created. Although, you can change these pairs, pairs of volumes that cannot be migrated will not be created.
- During migration, data can still be accessed.
- When data migration is completed normally, the LDEV ID and host I/O of the migration source volumes are changed to those of the migration target volumes. The procedure for accessing migrated data is not changed before or after migration.
- When you create a migration task, you can select whether to delete the migration source volume after the migration completes. As an alternative, after migrating, you can delete the data of the source volumes.
- After volumes are migrated to DP volumes, unused capacity in the target volumes is returned to the pool (zero page reclaim) if the zero page reclaim option is on. We recommend that the migration source volume be deleted if it is a DP volume and you will not use it. If you do not delete the

migration source volume, we recommend that the zero page of the migration source volume be discarded after the migration task finishes.
- You cannot work with volumes (such as copy pair or volume allocation) until the migration tasks are completed.

You must have a Tiered Storage Manager license registered to perform Data Migration.

Related tasks

- Migrating data for volume performance on page 459
- Migrating data to a different storage system on page 460

Related references

- Notes on performing data migration on page 452
- Conditions for data migration on page 453

Notes on performing data migration

Number of volumes

- The number of volumes that are candidates for migration target volumes are narrowed down when selecting a migration target volume. If you want to display all the volumes, in the `migration.properties` file on the Device Manager server set a value for the `migration.plan.candidateVolumeCountLimit` property.

  For details on the properties file, refer to the *Hitachi Command Suite Administrator Guide Administrator Guide*. To avoid insufficient memory, we recommend that you use the default.
- The maximum number of candidate volumes that can be displayed:
  - Migration target: 2,000
  - Migration source: 500

  Exceeding these values will produce an error.

System load

Data migration tasks take a long time from start of execution to completion. This might affect the performance of access from hosts to target volumes. To avoid executing such tasks while the access frequency is high, check the time required to complete the tasks and set a task execution schedule before registering the tasks. Actual required time might differ considerably in cases such as performing migration for the first time, migrating a volume smaller than 170 GB, or migrating many volumes at the same time. When storage systems or hosts are under a heavy load, migration might fail. In this case, reduce the load on the storage systems, and then perform data migration again.

Migration failures

- If a data migration fails, the integrity of data on the migration target cannot be ensured. Therefore, data migration must be restarted.
- If a data migration fails and operations for the migration target (such as re-execution of data migration and volume allocation) cannot be performed, the storage system must be refreshed.
- If the Hitachi Command Suite service stops during the execution or cancellation of a migration task, update the storage system information after restarting the service. You can use the **`hcmd64ssrv`** command or Hitachi Command Suite > Start - HCS (Windows). For more information, see the *Hitachi Command Suite Administrator Guide*.

Related concepts

- About data migration on page 451

Related tasks

- Migrating data for volume performance on page 459
- Migrating data to a different storage system on page 460

Conditions for data migration

Data migration conditions include conditions that apply to all volumes (volume migration requirements), and conditions that depend on the type of volume (volume migration requirements related to volume type).

For some volumes, when you try to select those volumes as migration-source volumes, the reasons why they cannot be selected as migration-source volumes are displayed on-screen.

Volume migration requirements

All volumes for which data migration is to be performed must meet the following conditions:

- Migration-source and migration-target volumes exist within one of the following storage systems (chassis):
  - VSP Gx00 models
  - VSP G1000
  - VSP G1500
  - VSP F1500
  - Virtual Storage Platform
  - Universal Storage Platform V/VM
  - HUS VM
- The emulation type, number of blocks, size, and T10 PI attribute are the same for the source and target migration volumes. However, note the following exceptions:
  - If the micro version of Universal Storage Platform V/VM is 60-03-20-00/00-xx or later and if both the migration source and the

migration target volumes are DP volumes, the emulation type can be between OPEN-V and OPEN-0V.

- If a volume whose emulation type is OPEN-V is a basic volume, and if the size of a migration target volume is larger than a migration source volume, a volume that has the same size of the migration source volume is re-created. The remaining capacity can be reused. If you set a value for the property `migration.plan.candidateCapacityGroupDisplayMaxCount` in the `migration.properties` file, you can specify how many volumes with a larger capacity than the migration source volume to display, in addition to the volumes with the same capacity as the migration source volume.

- The LDKC:CU:LDEV number for the volume is in the range from 00:00:00 to 00:FE:FF
- Volumes can be accessed, and are not reserved by other programs, Device Manager, or tasks.
- The emulation type is neither 3390-V nor 6588-A.
- The emulation type is not OPEN-V, and both the migration source and the migration target volumes are normal, or both are CVS volumes.

Volume migration requirements related to volume type

Whether you can select a volume as a migration source or migration target depends on the volume type. The following table shows which volume types can and cannot be selected as migration sources and targets.

Table 17 Migration sources and targets by volume type and state

Volume Type, State		Source	Target
Basic volumes		Y	Y
External volumes		Y	Y
LUSE volumes		Y[1]	Y
DP volumes[2]	HDT volumes to which a data placement profile is applied.	N	N
	Virtual volumes that are not associated with a DP pool.	N	N
	Volumes other than the above	Y	Y
DP pools		N	Y[3]
DP pool volumes		N	N
Volumes for which Cache Residency Manager is set		N	N
Command Device		N	N
Volumes that are configured as a Cross-system Copy pair		N	N
NAS system volumes		Y	N
System LUs in VSP G400, G600, G800, and VSP F400, F600, F800 storage systems with NAS modules		N	N
System Disk		N	N

Volume Type, State			Source	Target
Volumes that are configured for concurrent copying			N	N
On-demand device volumes			N	N
Intermediate volumes			N	N
Mainframe volumes for which a path is set			Y	N
Mainframe volumes for which SYSPLEXID is set			Y	Y
Mainframe volumes for which DEVN is set			Y	Y
Mainframe volumes for which VOLSER is set			Y	N
Source volumes for online data migration			Y	N
Volumes that are mapped by the nondisruptive migration attribute			N	N
Volumes that belong to the source group for online data migration			N	N
Quorum disks used for High Availability Manager			N	N
Volumes of ALU or SLU			N	N
External volumes used for data direct mapping			N	N
Volumes created when the drive type is FMC and volumes that belong to the parity group for which accelerated compression is enabled.			N	N
Deduplication System Data Volume			N	N
High Availability Manager	P-VOL S-VOL		N	N
ShadowImage, ShadowImage for Mainframe	Volumes for which pairs are configured with a P-VOL and three S-VOLs		N	N
	Volumes for which pairs are configured with an S-VOL and an SP-VOL		N	N
	Volumes for which pairs are configured with an SP-VOL and two S-VOLs		N	N
	Volumes other than the above		Y	N
TrueCopy Synchronous, TrueCopy for Mainframe	P-VOL S-VOL	PSUS PSUE [Suspended] PAIR [Duplex] (PAIR and Duplex can be a migration source when using TrueCopy Synchronous with VSP Gx00 models, VSP G1000, VSP G1500, VSP F1500, Virtual Storage Platform, Universal Storage Platform V/VM, or HUS VM)	Y[4] (In the case of an internal storage system)	N
		Volumes other than the above	N	N
Universal Replicator	P-VOL S-VOL	COPY PAIR [Pending duplex]	N	N

Volume Type, State			Source	Target
		[Duplex]		
		Volumes other than the above	Y[4]	N
	JNL-VOL		N	N
Copy-on-Write Snapshot	P-VOL	PAIR (1 to 63 generations)	Y	N
		Volumes other than the above	N	N
	V-VOL POOL		N	N
Thin Image	P-VOL		N	N
	V-VOL		N	N
	POOL		N	N
Data Retention Utility[5]	Read/Write		Y	Y
	Volumes other than the above		Y	N
Volumes that are configured for XRC	Suspended status		Y	N
	Volumes other than the above Volumes that are configured for XRC		N	N
Volumes for which virtual IDs are specified			Y	Y[6]
Volumes that are a global-active device pair.			N	N
Volumes reserved as global-active device secondary volumes.			N	N
Volumes that are quorum disks in a global-active device environment.			N	N
Compatible FlashCopy® V2, or Compatible FlashCopy®Compatible Software for IBM FlashCopy SE		Target volumes	N	N
		Volumes other than the target volumes	Y	N

Legend:
- Y: Can be selected
- N: Cannot be selected
- [ ]: Names indicated are those for mainframe-based programs

Notes:

1. When the migration source volume is a LUSE, a migration target must be selected for each volume that constitutes the LUSE. Migration target volumes must also satisfy the following conditions:
- All of the migration target volumes are internal volumes, or all are external volumes.
- The I/O suppression mode, cache mode, and CLPR of all migration target volumes are the same.
- Migration target volumes are not DP volumes.

Also, for HCS to automatically create pairs for migration sources and migration targets, we recommend that the following conditions also be met:
- If the migration target volumes are external volumes, all of the volumes are in the same external storage system.
- The RAID level, drive type, and SLPR (storage management logical partition) of all migration target volumes are the same.

Note that if you are migrating a LUSE volume of a Universal Storage Platform V/VM storage system for which the micro version is earlier than 60-05-12-00/00, you must not perform I/O operations on the volumes while migration is being executed.

Volume Type, State	Source	Target
2. The migration-source volumes and migration target volumes or pools are in different DP pools. **3.** The following conditions must be satisfied: • The status is not blocked. • The DP pool has more unused capacity than the migration-source volumes. • Used DP pool capacity does not exceed both used thresholds 1 and 2 (even after a migration is performed). • DP pool subscription does not exceed the subscription limit (even after a migration is performed). • When selecting migration-target resources, if the option to create new DP volumes is selected, the number of virtual volumes in the DP pool is less than the limit (VSP G800: 14,080; VSP G600/VSP G400: 4,096; VSP G200: 2,048; VSP G1000, VSP G1500, and VSP F1500: 63,232; Virtual Storage PlatformVirtual Storage Platform: 63,232; Universal Storage Platform V/VM: 8,192; HUS VM: 14,080). • If the migration-source volumes are DP volumes, the DP pools at the migration source and migration target must be different. • The Data Direct Mapped attribute must not be enabled. **4.** The CLPRs for the migration target and migration source are the same. **5.** If an external storage system is locked (not Read/Write), it cannot be a migration source or migration target. **6.** Volumes with virtual IDs specified cannot be selected as migration target volumes when: • The migration target is a DP pool, and Create new HDP volumes is selected in the migration target creation options. • The specified capacity differs from the capacity of the migration source.		

Reasons why volumes cannot be selected as migration sources

If a volume cannot be made a migration source and the reason for it is displayed, in the Data Migration wizard, a character string is displayed in the Migration Restriction column for Selected Column on the Migration Source page. The following table shows the character strings and corresponding reasons why a volume cannot be made a migration source.

Table 18 Volumes that cannot be used as migration sources

Displayed character string	Reason preventing use as a migration source
Array Type	The volume is part of an unsupported storage system.
Cache Residency Manager	Cache Residency Manager or Cache Residency Manager for Mainframe is set for the volume.
Command Device	The volume is used as a command device.
Copy-On-Write Snapshot Status	The volume is one of the following: • The volume is the P-VOL of a Copy-on-Write Snapshot pair, its status is PAIR, and the number of generations is 64 • The volume is the P-VOL of a Copy-on-Write Snapshot pair and its status is not PAIR

Displayed character string	Reason preventing use as a migration source
	• The volume is the V-VOL of a Copy-on-Write Snapshot pair • The volume is the POOL of a Copy-on-Write Snapshot pair
Deduplication System Data Volume	The volume is used as a system data volume for deduplication.
Dynamic Provisioning [Pool Volume]	The volume is a DP pool volume.
Dynamic Provisioning [Unassigned]	The volume is a virtual volume and is not associated with a DP pool.
External volume for which the Data Direct Mapping attribute is enabled	External volume used for data direct mapping
Externally Locked	The volume is locked by Data Retention Utility (not Read/Write).
GAD Reserved	Volumes reserved as global-active device secondary volumes.
global-active device	Volumes that are a global-active device pair.
Intermediate volume	The volume is an intermediate volume.
In the parity group, the volumes for which Accelerated Compression is enabled	The volume is in a parity group for which accelerated compression is enabled.
Migration Reserved	The volume has already been reserved as a migration target.
NAS Platform (System LU)	The volume is a system volume for the NAS module that is installed with VSP G400, G600, G800, and VSP F400, F600, F800.
Not Acquired Volume Information	Information for the volume cannot be acquired.
On Demand Device	The volume is an on-demand device volume.
Reserved	The volume has already been reserved in Device Manager.
Role	Only the View role is assigned.
ShadowImage Configuration [Leaf]	The volume is part of a ShadowImage or ShadowImage for Mainframe S-VOL, and paired with an SP-VOL.
ShadowImage Configuration [Node]	The volume is part of a ShadowImage or ShadowImage for Mainframe SP-VOL, and paired with two S-VOLs.
ShadowImage Configuration [Root]	The volume is part of a ShadowImage or ShadowImage for Mainframe P-VOL, and paired with three S-VOLs.
System Disk	The volume is a system disk.
Thin Image	This is a Thin Image P-VOL, V-VOL, or Pool.
TrueCopy Asynchronous Status	The volume is one of the following: • The volume is part of a TrueCopy Asynchronous pair, and the status is neither PSUS nor PSUE • The volume is part of a TrueCopy for Mainframe pair, and the status is not suspended
TrueCopy Synchronous Status	The volume is one of the following:

Displayed character string	Reason preventing use as a migration source
	• The volume is part of a TrueCopy Synchronous pair, and the status is other than PSUS, PSUE, or PAIR • The volume is part of a TrueCopy for Mainframe pair, and the status is not suspended
Universal Replicator Status	The volume is one of the following: • The volume is part of a Universal Replicator pair, and the status is either COPY or PAIR • The volume is part of a Universal Replicator for Mainframe pair, and the status is either Pending duplex or Duplex • The volume is part of a Universal Replicator journal volume • The volume is a Universal Replicator for Mainframe journal volume
Volume Migration V2	Volume Migration V2 is not installed in the storage system.

Related concepts

- About data migration on page 451

Related tasks

- Migrating data for volume performance on page 459
- Migrating data to a different storage system on page 460

Migrating data for volume performance

You can migrate data based on the results of monitoring volume performance.

If Hitachi Command Suite is linked with Hitachi Tuning Manager, you can display the Historical Report and check performance information by clicking **Analyze Performance**. For information about changing the report window settings for the Historical Report, see the *Hitachi Command Suite Tuning Manager User Guide*.

Before you begin

- Register a Tiered Storage Manager license.
- Set up an external connection to a VSP Gx00 models, VSP G1000, VSP G1500, VSP F1500, Virtual Storage Platform, Universal Storage Platform V/VM, or HUS VM storage system.
- Create a tier in which the target mainframe volumes are grouped (only when migrating mainframe volumes).
- Gather the following information:
  - Migration source host, file server, logical group, tier, or storage system.
  - Migration target DP pools, tier or migration target volumes.

Procedure

1. In the **Resources** or **Mobility** tab, select a migration source host or logical group. From a list of hosts or logical groups, you can also select migration source volumes. If migrating data from a mainframe volume, expand the **Tiers** tree on the **Resources** tab or the **Mobility** tab, and then select the corresponding tier.
2. Click **Migrate Data**.
3. Verify that the correct migration source volumes are selected, and then click **Next**.
4. Specify the migration target, DP pool, or tier and then click **Next**. When the parity group selection rule is set, a candidate volume that is appropriate as a migration target is selected depending on the application that is being used or the usage rate of the parity group. Specifically, the system attempts one of the following:
   - **Balance Capacity**: Selects a migration target candidate so that the remaining capacity of the targeted parity group is equalized with other parity groups.
   - **Maximum Coverage**: Selects a candidate for the migration target from as many parity groups as possible.
   - **Minimum Coverage**: Selects a candidate for the migration target from as few parity groups as possible.

   If Hitachi Command Suite is linked with Hitachi Tuning Manager, parity group busy rate for the migration target volume can be checked.
5. Confirm that all migration source volumes and target volumes or DP pools are paired properly, confirm that no other problems exist, and then click **Next**. If necessary, change the pairs. To change a pair, first select the volume pair, then from the **Other Candidates for Migration** list, select the target volume to change to, and then click **Map Volume**.
6. Specify plan options as necessary, and then run the task.
7. In the tasks list, confirm the result of the task. After a task is created, you must not perform any operations on volumes included in the task by using Hitachi Command Suite or other programs until the task has finished executing.

Related concepts

- About data migration on page 451

Related tasks

- Migrating data to a different storage system on page 460

Related references

- Conditions for data migration on page 453

Migrating data to a different storage system

You can migrate data to a different storage system.

If Hitachi Command Suite is linked with Hitachi Tuning Manager, you can display the Historical Report and check performance information by clicking **Analyze Performance**. For information about changing the report window settings for the Historical Report, see the *Hitachi Command Suite Tuning Manager User Guide*.

Before you begin

- Register a Tiered Storage Manager license
- Set up an external connection to the VSP Gx00 models, VSP G1000, VSP G1500, VSP F1500, Virtual Storage Platform, Universal Storage Platform V/VM, or HUS VM storage system.
- Create a tier in which the target mainframe volumes are grouped (only when migrating mainframe volumes).
- Identify the migration source host, file server, logical group, tier, or storage system.
- Identify the migration target DP pools, tier, or migration target volumes.

Procedure

1. On the **Resources** tab, select **Storage Systems**.
2. Expand the tree for the source storage system that is to be migrated, select the volumes to be migrated from the list of external storage systems, and then click **Migrate Data**.
3. Confirm that the correct migration source volumes are selected, and then click **Next**.
4. Specify the migration target, DP pools, or tier and then click **Next**. When the parity group selection rule is set, a candidate volume that is appropriate as a migration target will be selected depending on the application that is being used or the usage rate of the parity group. Specifically:
   - **Balance Capacity**: System attempts to select a migration target candidate so that the remaining capacity of the targeted parity group is equalized with other parity groups.
   - **Maximum Coverage**: Select a candidate for the migration target from as many parity groups as possible.
   - **Minimum Coverage**: Select a candidate for the migration target from as few parity groups as possible.

   If Hitachi Command Suite is linked with Hitachi Tuning Manager, parity group busy rate for the migration target volume can be checked.
5. Confirm that all migration source volumes and target volumes or DP pools are paired properly, confirm that no other problems exist, and then click **Next**. If necessary, change the pairs. To change a pair, first select the volume pair, then from the **Other Candidates for Migration Target** list, select the target volume to change to, and then click **Map Volume**.
6. Specify plan options as necessary and then execute the task.

7. In the tasks list, confirm the execution result of the task. After a task is created, you must not perform operations on volumes included in the task by using Hitachi Command Suite or other programs until the task has finished executing.

Related concepts

- About data migration on page 451

Related tasks

- Migrating data for volume performance on page 459

Related references

- Conditions for data migration on page 453

Migrate data dialog box

Data migration provides the capability for moving data from a set of source volumes to a set of destination volumes or DP pools, including those in externally-connected storage systems.

The Migrate data dialog box is a wizard that guides you in configuring the task for migrating selected data from a source location to a destination target.

For example, when you select a specific host, storage system, pool type, and click Apply, the corresponding volume information is displayed in the list of selected volumes because this information is already known and is non-configurable. To create a data migration task, you must complete a four-step configuration process using the migrate data wizard.

When you enter the minimum required information in a dialog box, review the Plan Summary, and click Submit to complete the configuration of this migration task. Click the Back button to modify the plan to meet your requirements.

The following table describes the dialog box fields, subfields, and field groups. A *field group* is a collection of fields that are related to a specific action or configuration. You can minimize and expand field groups by clicking the double arrow symbol (>>).

As you enter information in a dialog box, if the information is incorrect, errors that include a description of the problem appear at the top of the box.

Table 19 Migrate Data dialog box

Field	Subfield	Description
1. Migration Source		
Step 1 - **Specify the criteria to select the source volumes.**		

Field	Subfield	Description
Logical group	-	Click Select to display the **Select Logical Group** dialog box from which you can select from the available types. After selecting, click OK to continue.
Host	-	The Host drop-down list displays available hosts from which you can select.
Storage system	-	The Storage system drop-down list displays available storage systems from which you can select. **Note:** If the selected storage system supports or has corresponding pools, the Pools drop-down list is active.
Pool	-	The Pool drop-down list displays the available pools from which you can select.
>> Advanced Options	(See the following fields for details)	Click >> Advanced Options to display the Volume criteria options.
Volume Criteria	Category	Select an option from the Category drop-down list (for example, Storage System).
	Attribute	Select an option from the Attribute drop-down list (for example, Storage System Vendor).
	Attribute Condition	Select an option from the attribute condition drop-down list (for example, is).
	Value	Select an option from the Value drop-down list). Click + to add additional volume criteria condition entries, or click - to delete a specific volume criteria condition entry.
	Show entries matching	Select Any or All from the Show entries matching drop-down list. The setting you choose determines whether the source volumes match the use of any or all of the specified criteria.
	Clear	Click Clear to delete all existing volume criteria selections.
	Apply	Click Apply to define the selected volume criteria settings to select the migration source volume.
Step 2 - **Verify the selected volumes and click [Next].**		
	Next	From the list of selected volumes, verify that these are the source volumes you want to migrate, and click Next. **Note:** To observe details about a specific volume, click its name in the Volume column to display the Detailed Information dialog box for that volume.
2. Migration Target		
Specify a pool, tier, or volume criteria for the migration targets.		
Storage system	-	Identifies the selected storage system that is the migration source.
Migration-Target Selection	-	Depending upon the selected storage system, choose one of the following available migration target options: • Pool • Tier • Volume Criteria
	Pool	Click Pool and select the pool migration target from the drop-down list.

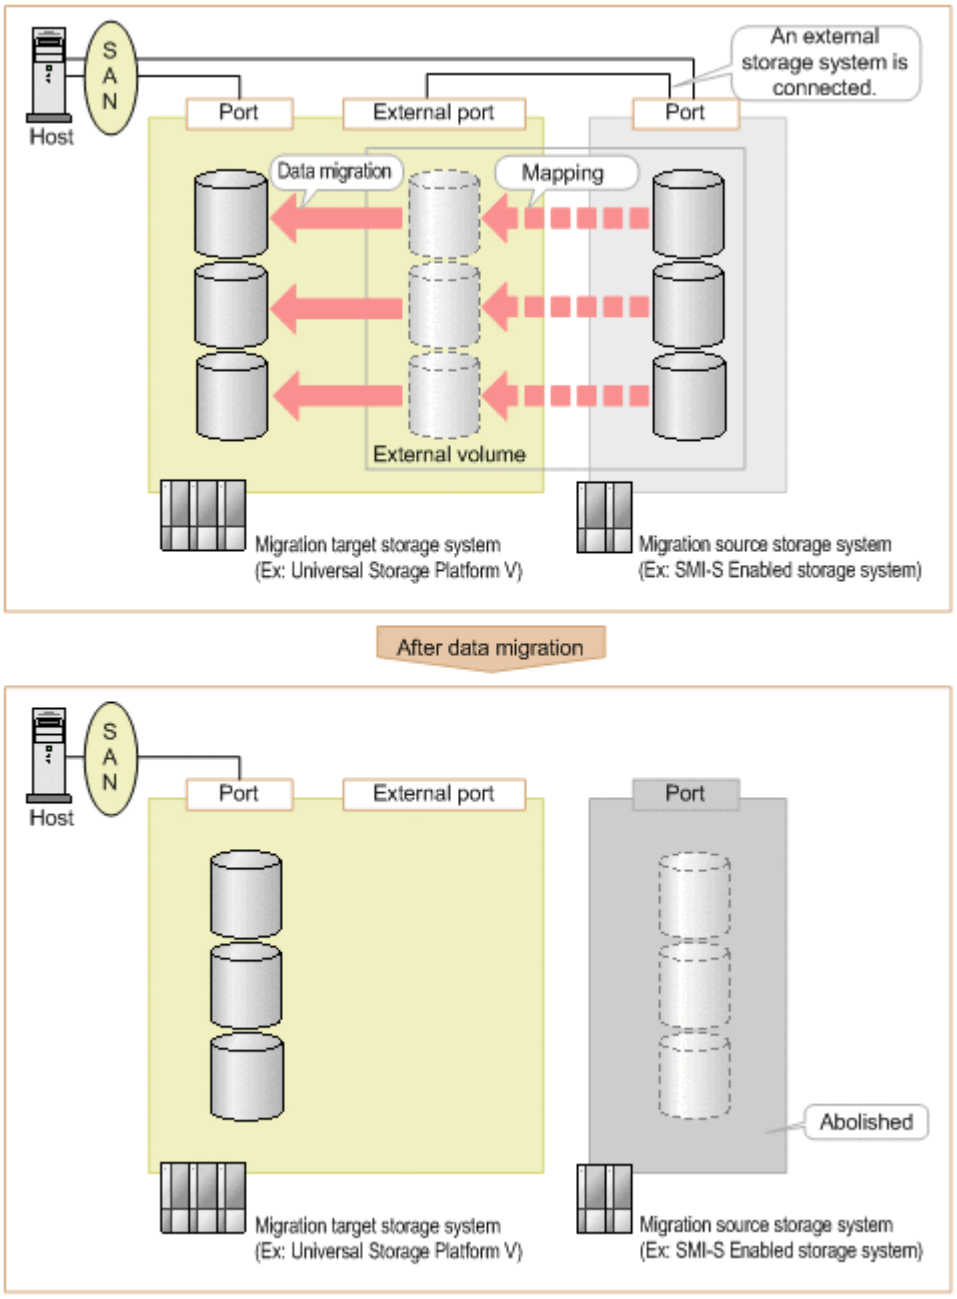
Field	Subfield	Description
Migration-Target Creation Options	-	Choose one of the following Migration-Target Creation Options options: • Create new DP volumes • Use existing DP volumes
	Create new DP volumes	**Note:** Volume criteria and Parity group selection rule are disabled when you select Pool and Create new DP volumes as your Migration Target Selection choices. If you selected Pool and this option, click Next to continue on to **Step 3** in this process.
	Use existing DP volumes	**Note:** Volume criteria is enabled and Parity group selection rule is disabled. If you select Pool, click Next and continue on to **Step 3** in this process.
>> Advanced Options	Migration-Target Selection: Pool	Click >> Advanced Options to display the Volume criteria options.
Volume Criteria	-	**Note:** The Volume criteria drop-down list choices you make affect each subsequent criteria option for the migration volume target. For details on setting these options, see >> Advanced Options in **Step 1** in this table.
Migration-Target Selection	-	Depending upon the selected storage system, choose one of the following available migration target options: • Pool • Tier • Volume Criteria
	Tier	Click Tier and select the tier from the drop-down list to be the tier migration target.
>> Advanced Options	Migration-Target Selection: Tier	Click >> Advanced Options to display the Volume criteria options. **Note:** The Volume criteria drop-down list choices you make affect each subsequent criteria option for the tier migration target. For details on setting these options, see >> Advanced Options in **Step 1** in this table.
	Parity group selection rule	Click Parity group selection rule, and select one of the following parity group selection rule options: • Balance Capacity (use the least utilized parity groups first) • Maximum Coverage (use the maximum number of parity groups) • Minimum Coverage (use the minimum number of parity groups) Click Next to conclude selecting the tier migration target, and continue on to **Step 3** in this process.
Migration-Target Selection	-	Depending upon the selected storage system, choose one of the following available migration target options: • Pool • Tier • Volume Criteria
	Volume criteria	Click Volume criteria and click >> Advanced Options to specify the volume criteria settings for the volume migration target.

Field	Subfield	Description
>> Advanced Options	Migration-Target Selection: Volume criteria	Click >> Advanced Options to display the Volume criteria options. **Note:** The Volume criteria drop-down list choices you make affect each subsequent criteria option for the migration volume target. For details on setting these options, see >> Advanced Options in **Step 1** in this table.
	Parity group selection rule	Click Parity group selection rule, and select one of the following parity group selection rule options: • Balance Capacity (use the least utilized parity groups first) • Maximum Coverage (use the maximum number of parity groups) • Minimum Coverage (use the minimum number of parity groups) Click Next to conclude selecting the migration target, and continue on to **Step 3** in this process.
3. Migration Pair		
Edit migration pairs.		
Analyze Performance	-	Click Analyze Performance to display a report of historic performance for one or more selected migration pairs (this requires a Tuning Manager license).
Do Not Migrate	-	Click Do Not Migrate to remove one or more selected migration pairs from migration consideration.
Restore Default Pairs	-	Click Restore Default Pairs to restore one or more selected migration pairs to default values.
Other Candidate pools for Migration	-	From the other candidate pools for migration list, click Map Pools to select other pools to map into the migration process. Click Next to conclude editing the migration pairs, and continue on to **Step 4** in this process. **Note:** If you selected any option besides Pool in Step 2 of this wizard, you would click Map Volumes to select other volumes to map into the migration process, and this list table would be titled Other Candidate volumes for Migration.
4. Show Plan and Option Setting		
Verify the plan for migrating data. You can also specify migration options		
Plan Summary	-	In the plan summary list, review the following migration option settings: • Storage System - listed system name and/or IP address • Migration Source - listed host name, logical group name and location, and pool name • Migration Target - listed pool or volume name • Estimated Time - listed time to complete migration in hh:mm:ss format If you disagree with the migration option settings, click Back and make the necessary changes.
>> Plan Details	-	Click >> Plan Details to display the migration pairs list. Review the listed migration pairs . If you disagree with the listed migration pairs, click Back and make the necessary changes.

Field	Subfield	Description
>> Migration Options	-	Click >> Migration Options to display the following migration options: • Shredding • Deleting • Zero page reclaim • Email address that received completion emails After choosing your option, enter a task name and description in the appropriate Task name and Description text boxes.
	Shredding	Click the Shredding option to enable the destruction (shredding) of deleted data from the source volume during the migration operation.
	Deleting	Click the Deleting option to enable deletion of the migration source volume after migration.
	Zero page reclaim	Click the Zero page reclaim option to enable the process where when volumes are migrated to DP volumes, any unused physical capacity in the target volumes is returned to the pool. Pools are examined by the firmware, which if it verifies that no data other than zeros are detected and physical storage is unmapped, it can be returned as free storage capacity.
	Email address that received completion emails	Click the Email address that received completion emails option and enter a valid email address into the adjacent text box.
>> Schedule	-	Schedule the running of the migration task by using one of the following options: • Now • Later **Note:** After you leave this wizard, you can click the Tasks & Alerts tab and view the status of any scheduled task.
	Now	Click Now to schedule the running of this migration task now (which is the default).
	Later	Click Later and specify a date and time to run the migration task in the Date and Time text boxes.
	View task status	If you select the View task status check box, you can check the task status in the tasks list in the Tasks & Alerts tab.
Submit	-	Click Submit to submit the migration task to the system.

Prerequisites for migrating data from an SMI-S enabled storage system to another storage system

You can migrate data between an old and new storage system for abolishing or merging systems by using the external storage connection function (Universal Volume Manager). The following figures provide an example of migrating data to Universal Storage Platform V when the SMI-S enabled storage system is abolished. To perform migration, use the GUI or CLI to migrate data by the system.



Related tasks

- Migrating data from an SMI-S enabled storage system to another storage system on page 467

Migrating data from an SMI-S enabled storage system to another storage system

You can migrate data from an SMI-S enabled storage system to another storage system.

Before you begin

- Register a Tiered Storage Manager license.
- Set up Tiered Storage Manager CLI (if necessary).
- Set up Device Manager CLI (if necessary)
- Connect the port on the SMI-S enabled storage system to the External port of the storage system by using Fibre Channel or Fibre Channel over Ethernet.
- Assign the volumes on the SMI-S enabled storage system to an External port (performed by the management tool of the SMI-S enabled storage system).
- Gather the following information:
  - Name or IP address of the allocation target host.
  - WWN and LUN associated with the migration source SMI-S enabled storage system that was assigned to the External port.

For details about the specified values specified for parameters and syntax of each CLI command of Device Manager and Tiered Storage Manager to be used in tasks, see the *Hitachi Command Suite CLI Reference Guide* and the *Tiered Storage Manager CLI Reference Guide*.

Procedure

1. Register the migration target storage system and the migration source SMI-S enabled storage system in Hitachi Command Suite. Use the GUI to register the storage systems, or run the **`AddStorageArray`** command (Device Manager CLI).

   Note: Registering or updating a storage system might take a long time depending on its configuration or scale, so you need to perform these operations systematically.

   If the migration target or source storage system has already been registered, you need to update the migration target or source storage system in order to obtain the latest configuration information.

2. Stop the hosts to which the storage systems are connected.

   Observe that the following steps (3 - 4) need to be performed while the host is stopped, so you need a plan that considers host operations.
3. Map the volumes of the migration source SMI-S enabled storage system to the migration target storage system as external volumes.

   Use the GUI to map the volumes or run the **`AddExternalArrayGroup`** command (Device Manager CLI).
4. Allocate the external volumes mapped to the migration target storage system to the host.

   Use the GUI to allocate the volumes, or run the **`AddLun`** command (Device Manager CLI).

5. Start the host, and then change the target that the host (applications) accesses from the migration source storage system to the migration target storage system.

   Perform this step by using the configuration tools of the host while applications are stopped.

6. Migrate data in the external volumes (migration source) to an internal volume of the migration target storage system (data migration). Use the GUI to migrate data, or run the following commands (Tiered Storage Manager CLI):
   - **`CreateStorageTier`** command
   - **`CreateMigrationGroup`** command
   - **`CreateMigrationPlan`** command
   - **`CreateMigrationTask`** command
   - **`ExecuteTask`** command
7. Unmap the external volumes (migration source) from the migration target storage system.

   Use the GUI to unmap (unvirtualize) the volumes or run the **`DeleteExternalArrayGroup`** command (Device Manager CLI).
8. Delete the migration source SMI-S enabled storage system from the management targets of Hitachi Command Suite.

   Use the GUI to delete a storage system, or run the **`DeleteStorageArray`** command in the Device Manager CLI.

Result

The data of the migration source SMI-S enabled storage system can now be used as an internal volume of the migration target storage system. The migration source SMI-S enabled storage system can be removed.

Related references

- Operations available to SMI-S enabled storage systems on page 65

12

Grouping resources

This module describes how to group logical groups and storage tiers (resources).

- Managing logical groups
- Managing storage tiers

Managing logical groups

This module describes how to manage the two supported logical group types: public and private. Public and private logical groups are listed in the navigation pane (Resources > Logical Groups) under the Public Logical Groups root folder and the Private Logical Groups root folder, respectively.

About logical groups

A *logical group* is a mechanism for the grouping of storage resources (hosts and volumes) by storage administrators, and can be organized according to business operations, geographic locations, or other organizational divisions.

Using grouped storage resources as logical groups allows you to:

- Create summary reports that indicate current resource use in a storage environment by organization.
  You can create logical groups that generate usage summary reports for each organization that uses storage, if you manage an entire system.
- Monitor volumes and perform storage system operations by grouping together managed resources (hosts and volumes).
  You can categorize volumes according to the organizations that use the storage, if you manage a part of a system. Grouping and categorizing volumes makes finding volumes easier and allows you to perform daily operations for volume resources.

You can add multiple hosts or volumes to a logical group at one time by specifying a range of IP addresses, or by specifying multiple items, such as operating system type and drive type. The same hosts or volumes can be registered to multiple logical groups.

When you need to, you can add or remove hosts and volumes from a logical group, or change the hierarchy, or structure, of existing logical group folders, depending on the logical group type. Any volumes allocated to a host in a logical group are automatically assigned to the logical group. Note that child groups cannot be created below host or volume groups (child groups can be created only below folders).

Hitachi Command Suite supports two types of logical groups: public and private. A key distinction between public and private logical groups is the root location in the work area for the user, and these two logical group types have the following basic structure in the navigation pane:

Public logical groups

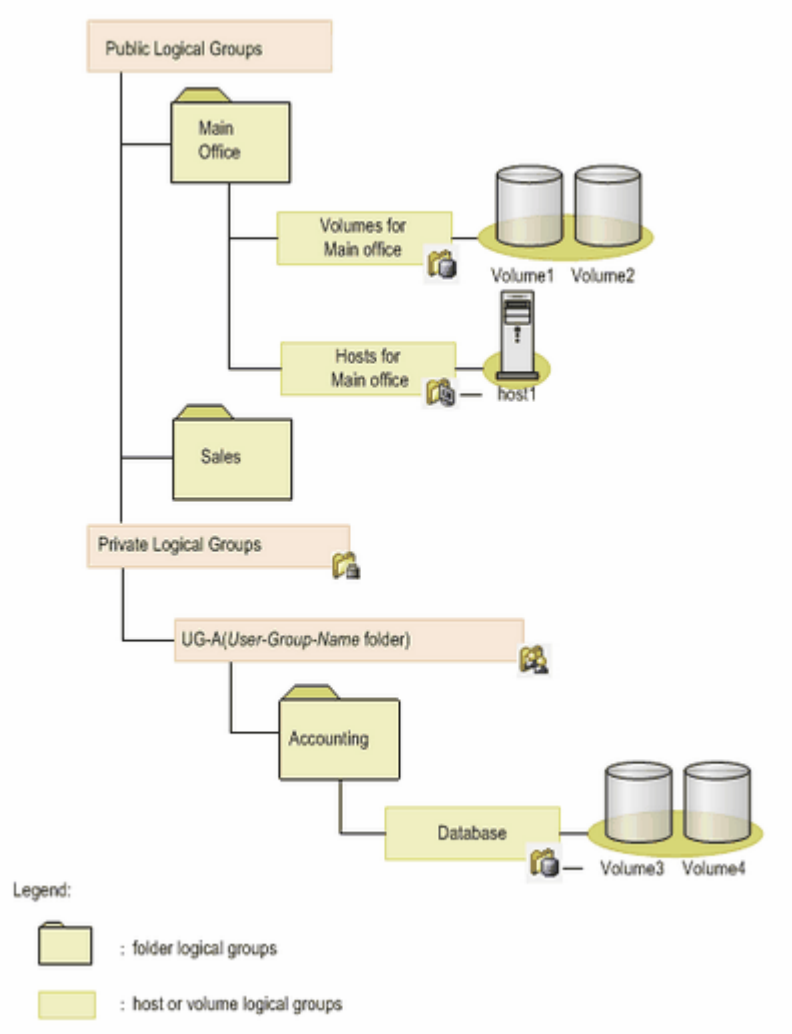
- Logical group folders that contains groups of hosts and volumes that are grouped by business function, geographic, or other organizational divisions.
- Logical groups of hosts.
- Logical groups of volumes.

Private logical groups

Top-level folder that corresponds to each user group.

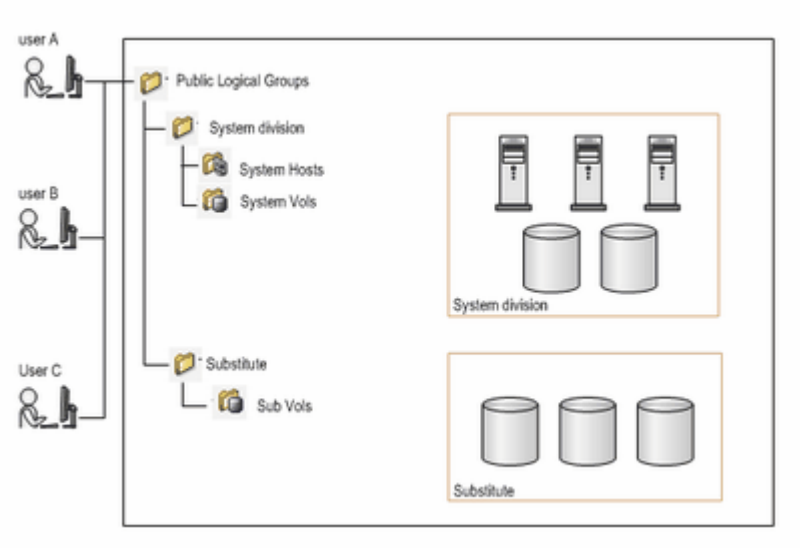
- Private logical group folder that contains groups of hosts and volumes that are grouped by business function, geographic, or other organizational divisions.
- Logical groups of hosts.
- Logical groups of volumes.

The following figure illustrates the structure of logical groups.



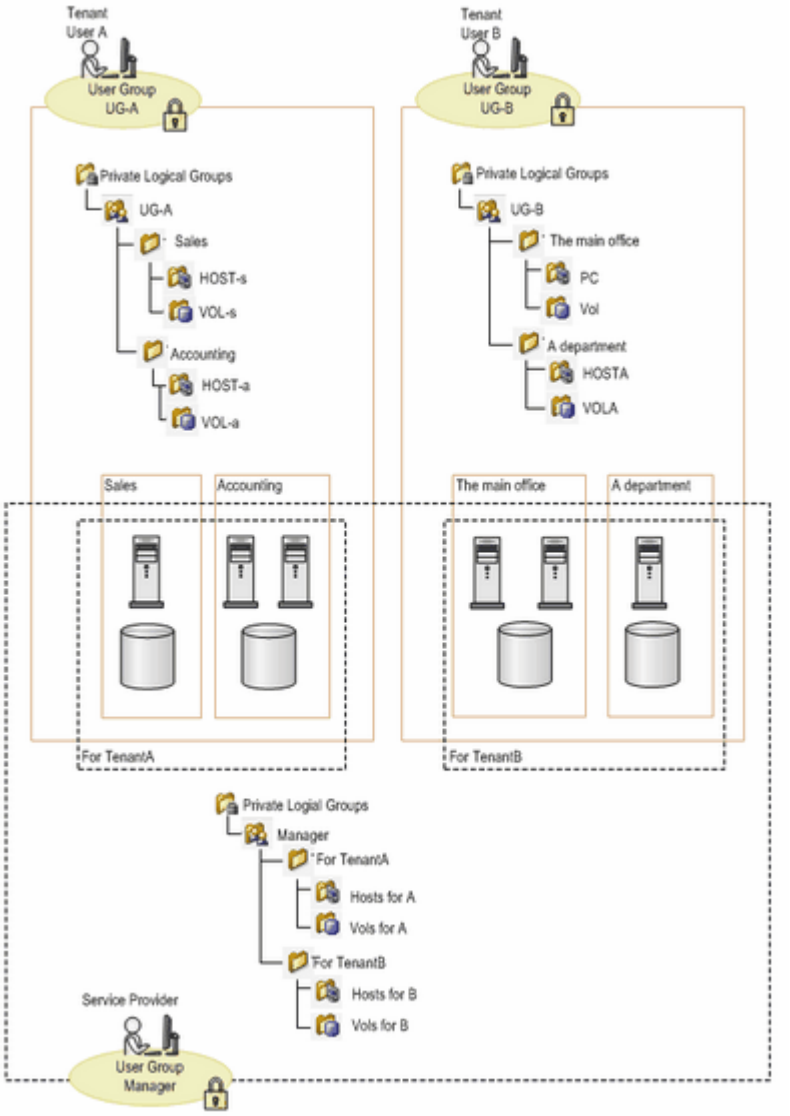
The following descriptions and figures show key differences between public logical groups and private logical groups.

- In public logical groups, any HCS user can access public logical groups.



- In private logical groups, only users that belong to specific user groups can access the private logical groups that are associated with the user groups to which they belong.
  The Private Logical Groups root folder can include separate private logical groups thereby maintaining a logical separation between business organizations or tenants, to provide for the secure sharing of resources.

For each of these tenant user groups, only users that belong to the user group can see or access the private logical groups for that user group.



When a user group is created, HCS automatically creates a top-level folder for each user group under the Private Logical Group folder. The default name of the folder is the user group name. All members of the user group can use this top-level folder as their work folder. In this work folder, users can create, edit, or delete private logical groups. You can change the default name for the top-level folder name.

However, if you edit the user group name, the names are not synchronized. When a user group is deleted, the top-level folder under the Private Logical Groups root folder is also deleted. Be aware that all folders under the top-level folder are also deleted when the user group is deleted.

The following table describes in more detail the structure for public and private logical groups and the access rules for each logical group type.

Table 20 Logical groups structure

Type	Description
Public Logical Group	Root folder for public logical groups
>*public-logical-group*	Each public logical group can be a logical group of hosts, a logical group of volumes, or a logical group folder that contains both hosts and volumes. All users can access these public logical groups.
Private Logical Group	Root folder for private logical groups
> *user-group-name* top-level folder	A top-level folder for each user group. This folder is created by HCS under the Private Logical Groups folder and can be used as a work folder for a user group. The default name of this top-level folder is the name of the user group. You cannot add, move, or delete these work folders except by adding or deleting a user group. When you delete a user group, the corresponding top-level folder and all included private logical groups under it are also deleted. Users that belong to the corresponding user group can create, edit, or delete private logical groups inside this folder.
>> *private-logical-group*	Each private logical group within a *user-group-name* folder can be a logical group of hosts, a logical group of volumes, or a logical group folder that contains both hosts and volumes. These logical groups can only be accessed by members of the user group to which this logical group belongs. Users that belong to the user group can create, edit, or delete private logical groups inside the *user-group-name* folder.

If you have a registered Tiered Storage Manager license, use the Mobility tab to check the volume status for each logical group or to optimize data arrangements. To efficiently find logical groups that have usage problems, use the Mobility tab to check their rankings. For example, such as any logical groups with a small amount of unused capacity or that use a large amount of capacity by drive type (for example: FMD, FMC, SSD, SAS, or SATA). The Tiered Storage Manager provides more detailed volume information such as IOPS, response time, and used capacity for each tier in the system in addition to the status information displayed by the Resources tab.

If you do not have a registered Tiered Storage Manager license, use the Resources tab to check the volume status for each logical group. For example, the Resources tab displays the following volume status-related categories displayed in the application pane:

- **Volume**
- **Storage System**
- **WWN Nickname**
- **Port**
- **Port Type**
- **LUN**
- **Capacity**

Related tasks

- Creating logical groups on page 477
- Editing logical groups on page 479
- Viewing logical group reports on page 478
- Performing operations from a logical group on page 479
- Deleting logical groups on page 480
- Creating user groups on page 134

Creating logical groups

HCS supports two types of logical groups: Public Logical Groups and Private Logical Groups.

Before you begin

- Verify that any host to be assigned to a logical group is a normal host, a virtualization server, or a virtual machine registered as a host.

Observe the following restrictions:

- File servers or NAS modules cannot be assigned to a logical group.
- Mainframe volumes (MF-Vols) cannot be assigned to a logical group.
- Child groups cannot be created below host or volume groups. Child groups can be created only below folders.

Procedure

1. On the **Resources** tab, select **Logical Groups**.
2. Create a logical group using one of the following methods:
   - For a Public Logical Group, select the Public Logical Groups root folder and click **Create Logical Group**.
   - For a Private Logical Group, select the Private Logical Groups root folder, select a top-level folder for each user group under the Private Logical Group folder, and then click **Create Logical Group**.
3. In the **Create Logical Group** dialog box, enter information in the **Name** and **Description** fields, and select one of the following **Create as** options: **Folder**, **Group of hosts**, or **Group of volumes**.
   - To create a folder, click the **Folder** option and click **OK**.
   - To create a group of hosts, click the **Group of hosts** option, and choose how to specify hosts for a logical group by selecting one of the following **Specify hosts** options: **By Search Criteria** or **Manually**. If you selected **By Search Criteria**, click **Add Criteria** in the list of criteria. In the **Add Conditions** dialog box, specify values for **IP Address Criteria**, **Name Criteria**, and **Operating System Criteria** as needed. You can also choose to click **Search Result Preview**, select one or more hosts in the list of hosts, and click **Add** to add them to the list of hosts, and then click **OK**.

If you selected **Manually**, click **Add Hosts** and in the **Add Hosts** dialog box, select one or more hosts from the lists of hosts, click **Add** to add them to the list of selected hosts, and then click **OK**.

- To create a group of volumes, click the **Group of volumes** option, and choose how to specify volumes for a logical group by selecting one of the following **Specify volumes** options: **By Search Criteria** or **Manually**.

  If you selected **By Search Criteria**, click **Add Criteria** in the list of criteria. In the **Add Conditions** dialog box, specify values for **Usage Criteria**, **Storage System Criteria**, and **Volume Criteria** as needed. You can also choose to click **Search Result Preview**, select one or more volumes in the list of selected volumes, click **Add**, and then click **OK**.

  If you selected **Manually**, click **Add Volume** and in the **Add Volumes** dialog box, select one or more volumes from the list of available volumes, click **Add** to add them to the list of selected volumes, and then click **OK**.

4. Repeat this process for each new logical group that you want to create.
5. To create a hierarchy of logical groups, repeat step 2.

Result

New logical groups appear on the Resources tab under Logical Groups.

Tip: To delete logical groups, select the logical groups to be deleted, click Delete Logical Groups, and click OK.

Related concepts

- About logical groups on page 472

Related tasks

- Editing logical groups on page 479

Viewing logical group reports

You can view summary reports for specified logical groups.

Procedure

1. On the **Resources** tab, click **Logical Groups**.
2. Expand the tree and select the target logical group.

Related concepts

- About logical groups on page 472

Related tasks

- Creating logical groups on page 477

Editing logical groups

You can edit a logical group to change the name and description, move a logical group to another logical group folder, select a new parent logical group, or edit a logical group to add or remove hosts or volumes.

If data placement profiles for HDT volumes are created, a change of logical groups is reflected for the structure of data placement profiles.

Procedure

1. On the **Resources** tab, click **Logical Groups**.
2. Select **Public Logical Groups** or **Private Logical Groups**.
3. In the list of logical groups on the application pane, click the row for the logical group you want to edit, and click **Edit Logical Group**.
4. Edit the name or description of the logical group.
5. Click **Change Location** to move the logical group to another logical group folder, select a new parent logical group, and click **OK.**
6. Select the logical group for the volumes or hosts that you want to add or remove, and click **OK** .
7. Click **OK** to finish editing the logical group.

Result

Changes are reflected in the Logical Groups tree in the navigation pane, and in the details of the logical groups.

Related concepts

- About logical groups on page 472

Related tasks

- Creating logical groups on page 477

Performing operations from a logical group

You can perform tasks on storage systems by starting from a logical group.

Procedure

1. On the **Resources** tab, select **Logical Groups**, and select either **Public Logical Groups** or **Private Logical Groups**.
2. Expand the appropriate logical groups tree and select the target logical group.
3. Select a host or volume.
4. Click on the task type you want to perform such as **Allocate Like Volumes**, **Edit Labels**, **Edit LUN Paths**, or **Export to CSV**, or click **More Actions** for additional choices.
5. If needed, update the task name and provide a description.

6. Click **Submit**.
7. You can check the progress and result of the selected task on the **Tasks & Alerts** tab, and verify the results.

Related concepts

- About logical groups on page 472

Related tasks

- Creating logical groups on page 477

Deleting logical groups

You can delete logical groups when they are no longer required.

Procedure

1. On the **Resources** tab, select **Logical Groups**.
2. Depending on the type of logical group you want to delete (public or private), do one of the following:
   - Click **Public Logical Groups** in the navigation pane, and then select a logical group from the list of logical groups in the application pane.
   - Expand **Private Logical Groups** in the navigation pane, select a logical group folder from the navigation pane, and then select a logical group from the list of logical groups in the application pane.
3. Click **Delete Logical Groups**.

> **Tip:** Deleting a user group deletes the private logical group folder and all the logical groups in that folder. You cannot delete a top-level folder for each user group under the Private Logical Group folder.

4. Click **OK** to confirm and delete the logical group.

Result

The logical group you deleted no longer appears in the list of logical groups.

Related tasks

- Deleting user groups on page 136

Managing storage tiers

This module describes how to manage tiers.

About tier-based storage management

In Hitachi Command Suite, you must have a Tiered Storage Manager license to perform tier-based storage management operations.

In a tier-based storage management system, storage system resources are grouped and managed in tiers. Tiers can be created for volumes or for Dynamic Provisioning (DP) pools to allow high-capacity resources to be centrally managed.

Tiers allow you to allocate data according to its purpose. This allocation can be done by grouping drives in a storage system by type, or by creating groupings of particular tasks or functions. For example, you can allocate frequently accessed data to a tier of high-reliability drives or storage systems.

You can also use storage system resources more efficiently by storing data that is only for archive purposes on a tier of SATA drives or older storage systems. A volume can also have a volume belong to multiple tiers.

When you are managing a large variety of resources in tiers, you can check the performance status of all the resources in the tiers at the same time, and you can select the migration-source and migration-target volumes for data migrations from the tiers. By selecting tiers that match all the requirements when volumes will be allocated to hosts, you can allocate all the necessary volumes from the selected tiers at the same time.

Related tasks

- Creating tiers on page 481
- Expanding a tier on page 482

Values to ensure acceptable tier performance

The values below ensure the operation of Tiered Storage Manager. Performance might slow down or memory might become insufficient if a value exceeds these values.

Table 21 Tiered Storage Manager operational values

Item	Value for which performance is ensured	Recommended value
Number of tiers in a single storage system	100	30 or less
Number of unused volumes a tier	1,300	1,000 or less

Creating tiers

Create tiers to prioritize data based on accessibility requirements.

Before you begin

- You must have a valid license for Tiered Storage Manager.
- Register one or more of the following storage systems:
  - VSP Gx00 models
  - VSP G1000
  - VSP G1500
  - VSP F1500
  - Virtual Storage Platform
  - Universal Storage Platform V/VM
  - HUS VM

Procedure

1. On the **Resources** tab, select **Tiers**.
2. Select the storage system for which to create a tier.
3. Click **Create Tier**.
4. Specify the appropriate information.
5. Click **OK**.

Result

The tier is created and added to the storage system.

Related concepts

- About tier-based storage management on page 481

Related tasks

- Expanding a tier on page 482
- Deleting tiers on page 483

Expanding a tier

Edit a tier to expand it.

Before you begin

A Tiered Storage Manager license must be registered.

Procedure

1. On the **Resources** tab, select **Tiers**.
2. Expand the tree, and select the storage system with the tier you want to expand.
3. From the list of tiers in the application pane, select the tier you want to expand and click **Edit Tier**.
4. Specify the appropriate values.
5. Click **OK**.

6. Confirm that the tier has expanded.

Related concepts

- About tier-based storage management on page 481

Related tasks

- Creating tiers on page 481

Deleting tiers

You can delete storage tiers that you no longer need.

Procedure

1. On the **Resources** tab, select **Tiers**.
2. Select the tiers to delete.
3. Click **Delete Tier**.
4. Click **Delete**.
5. Confirm that the tier is deleted.

Result

The tiers you deleted no longer appear in the list of tiers.

13

Managing storage system information

This module describes how to manage resource labels and WWN nicknames, how to search managed resources, and how to generate storage system reports.

- Managing resource labels
- Managing WWNs by using nicknames
- Searching HCS resources
- Generating resource reports
- Using reports to verify system changes

Managing resource labels

This module describes how to manage resource labels.

About managing resource labels

User-defined resource labels can be assigned to volumes and DP pools. Specifically, resource labels can be volume labels, or DP pool names.

Resource labels help you:

- Search for volumes or DP pools
- Differentiate between listed volumes or DP pools
- Confirm the correct resource is being acted upon

When searching resource labels, you can:

- Allocate volumes or DP pools
- Edit volume LUN paths

Resource labels are assigned using:

- Device Manager
- Storage Navigator

For storage systems where resource labels have already been configured using Storage Navigator, you can apply (import) the resource labels to Device Manager in a batch operation.

If Device Manager and Storage Navigator are both actively used, and HCS is the primary management software, labels assigned in Device Manager can be applied (exported) to storage systems managed by Storage Navigator. Volume labels can be assigned when creating or allocating volumes depending on their use for hosts or applications. For VSP Fx00 models, VSP Gx00 models or VSP G1000, VSP G1500, or VSP F1500, importing and exporting labels is unnecessary because the label information from the storage system and the HCS labels will always match.



Note: Due to differences between Device Manager and storage systems, incompatible character type or string length issues will be reported with pop-up messages. Use the message to resolve resource label naming issues.



Tip: You can apply any current label to the storage system by specifying the following settings:

- For volume labels, select Reflect current labels to the storage systems.
- For DP pool names, select Reflect this pool name to the storage system.

Related tasks

- Editing resource labels on page 487
- Searching resource labels on page 487

- Importing storage system resource labels on page 488

Editing resource labels

You can create or edit resource labels for volumes or DP pool names.

Procedure

1. On the **Resources** tab, expand **Storage Systems** to display the list of volumes or DP pools.
2. Select a volume or DP pool.
3. Choose one of the following options:
   - Click **Edit Labels** to edit a volume label.
   - Click **Edit Pool** to edit a DP pool name.
4. Specify a new label or DP pool name.
5. (Optional) Update the task name and provide a description.
6. (Optional) Expand **Schedule** to specify the task schedule.

   You can schedule the task to run immediately or later. The default setting is **Now**. If the task is scheduled to run immediately, you can select **View task status** to monitor the task after it is submitted.
7. Click **Submit.**
   If the task is scheduled to run immediately, the process begins.
8. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.

Result

The revised volume or DP pool name appears in the list of volumes or DP pools.

Related concepts

- About managing resource labels on page 486

Searching resource labels

You can search, filter and sort on resource labels to locate volumes or DP pools and perform actions on them.

Procedure

1. Locate volumes or DP pools by using one of the following methods:
   - Using the search box, enter a volume label or DP pool name, and perform a search. You can use **Search Keyword for** to specify the resources you want to search, for example **Volumes**.
   - From the search box, you can also search volume labels by using **More Searches**, selecting Volume, and selecting the category Label from the Advanced tab.

- To filter, from the **Resources** tab expand the tree to display the list of volumes or DP pools, and click **Filter** to specify the volume label or DP pool name.
- To sort, from the **Resources** tab expand the tree to display the list of volumes or DP pools, and click **Label** or **Name** column headings to sort by volume label or DP pool name.

2. Select volumes or DP pools from the search, filter, or sort results.
3. Perform the desired task.

Related concepts

- About managing resource labels on page 486

Importing storage system resource labels

In an environment where storage system resource labels have already been set using Storage Navigator, you can apply (import) the storage system resource labels to Device Manager.

Note: Applying storage system resource labels overwrites all the volume labels and DP pool names in HCS for the selected storage system. For VSP Fx00 models, VSP Gx00 models or VSP G1000, VSP G1500, or VSP F1500, refreshing the labels is unnecessary because the label information from the storage system to HCS will always match.

Procedure

1. On the **Resources** tab, from the list of **Storage Systems** in the application pane, select the storage system whose labels you want to apply.
2. Click **Refresh Labels**.
3. Select the resource labels to apply, and submit the task.
4. Check the task results in **Tasks & Alerts**.

Result

Make sure the labels are changed in the volume list, and that pool names are changed in the DP pool list.

Related concepts

- About managing resource labels on page 486

Managing WWNs by using nicknames

This module describes how to specify, change or delete WWN nicknames.

About managing WWNs by using nicknames

Each host bus adapter (HBA) has a unique World Wide Name (WWN) identifier that consists of a 64-bit physical address. The purpose of specifying

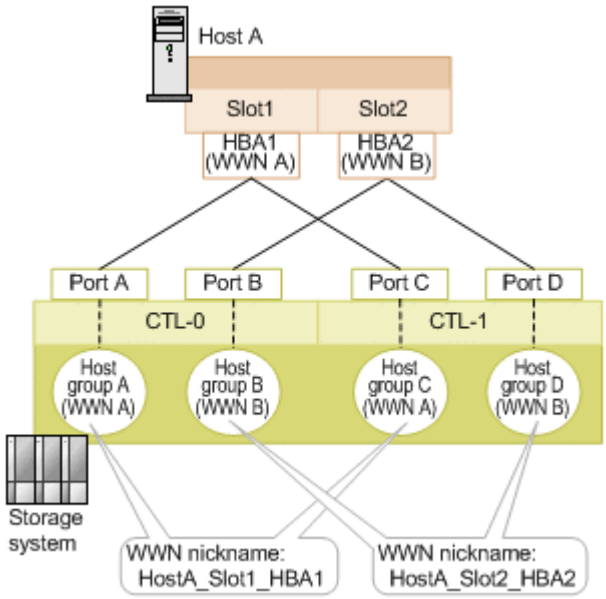
WWN nicknames is to easily identify accessible HBAs that are registered in a host group.

If one HBA WWN belongs to multiple host groups, you can set a nickname in the WWN information that is defined for each host group. WWN nicknames that are specified for a host group can provide information about the I/O path (host, host port, fabric, or storage port) of volumes that are linked to the host group.

During configuration of SAN fabrics and paths, the HBA WWN nickname helps you confirm physical configuration information. You can:

- Check the storage HBA WWN nicknames in the list of volumes.
- Specify or change the storage HBA WWN nicknames related to the host.
- View the HBA WWN nickname to confirm that the host port of the path to which a volume will be allocated is appropriate.
- Change the storage configuration, such as replacing an HBA and specifying the nickname for the replacement storage HBA WWN that is registered in the host group.

The following figure illustrates the WWN nicknames created using the host name, slot number, and HBA port.



WWN nickname examples:

- (Host Name)_(Slot Number)_(HBA Port)
- (Host Name)_(Slot Number)_(Fabric Number)
- (Host Name)_(Slot Number)_(Fabric Number)_(Port Number)

There are two options for editing a WWN nickname that is defined in a host group:

- Collectively edit the WWN nickname for all related host groups. When this option is selected, the WWN nickname is applied to all host groups that have the same WWN.
- Edit the WWN nickname of an individual host group. When this option is selected, only the WWN nickname of the individual host group is changed.

Tip: If a nickname is set to a WWN by using a storage system management tool, such as Storage Navigator, you can view the latest WWN nickname by refreshing the storage system information. Also, if there are multiple WWN nicknames set to the host groups on the storage system, only one WWN nickname for the port on the host is selected and displayed automatically.

Related tasks

- Editing a WWN nickname on page 490

Editing a WWN nickname

You can specify a new WWN nickname, or change or delete an existing nickname, for storage HBA WWNs that are registered in a host group.

Before you begin

Register HBA WWNs for which nicknames will be specified in the host group.

Procedure

1. On the **Resources** tab, select **Hosts**.
2. After selecting the target operating system, select a target host, and then click **More Actions > Edit WWN Nicknames.**
3. From the **Storage system** drop-down list, select a storage system.
   - **All Storage Systems**: Displays a list of WWNs that belong to the host group related to the selected host.
   - Specific storage system: Displays a list of WWNs that belong to the host group related to the selected host and storage system.
4. From **Edit option**, select one of the following:
   - **Edit WWN nickname for all related host groups:** Applies the WWN nickname setting to all storage host groups.
   - **Edit WWN nickname for individual host groups:** Applies the WWN nickname to individual host groups in the storage system selected.
5. In the **Nickname** column, enter a WWN nickname or modify the existing WWN nickname. If **All Storage Systems** and **Edit WWN nickname for all related host groups** is selected, the minimum number of character limit allowed is applied to all storage systems.

Tip: To delete a WWN nickname, clear the text.

6. (Optional) Update the task name and provide a description.
7. (Optional) Expand **Schedule** to specify the task schedule.
   You can schedule the task to run immediately or later. The default setting is **Now**. If the task is scheduled to run immediately, you can select **View task status** to monitor the task after it is submitted.
8. Click **Submit.**
   If the task is scheduled to run immediately, the process begins.
9. (Optional) Check the progress and result of the task on the **Tasks & Alerts** tab. Click the task name to view details of the task.

Result

The WWN nickname edits are applied to the WWNs registered in the host group.

Related concepts

- About managing WWNs by using nicknames on page 488

Searching HCS resources

This module describes using search to locate, work with, and report on storage resources.

About searching HCS resources

Search finds storage and host resources quickly and easily, and allows you to work with the resources from the search results.

Search is useful for:

- Investigating and working with storage resources - for example, you can quickly find a host and allocate a volume to the host. Note that when you click on a search results link, you will automatically go to the resource in the Resources tab so you can work with the resource, or you will get more detail on the item, as with volumes.
- Reporting on storage resources - for reporting, search results and related detail can be saved as CSV files using the Export to CSV button, or from the Actions menu. Note that CSV (comma separated values) files can be opened by many spreadsheet applications for review or editing.

Search methods include:

- Keyword search - a search on user entered text. You can enter a full or partial text string and search specific fields related to storage resources, such as:
  - Storage Systems - search names, IP addresses, and serial numbers
  - Hosts - search names, IP addresses, iSCSI names, and WWNs
  - Volumes - search volume IDs that match the keyword, or volume labels
  - Parity Groups - search parity group names
  - DP Pools - search IDs that match the keyword, and names

    - Logical Groups - search names or descriptions
- Criteria search - a search on attribute information, for example a volume search where the volume belongs to a RAID 5 parity group.
    - Criteria searches can be saved by the user, and optionally, shared with others
    - Note that saving searches requires a Tiered Storage Manager license

Related tasks

- Keyword search on page 492
- Criteria search on page 493

Keyword search

To perform a keyword search, locate the Search box on the upper right side of the GUI. The search box is always visible.

Procedure

1. Enter a full or partial text string in the **Search** box. A partial string will generally produce more search results.
2. By default, all resources are searched when you press Enter. You can narrow the scope of your search using the **down arrow**, selecting **Search Keyword for**, and selecting a specific resource to search on. This executes the search and displays results. Each resource has an icon that appears in the search box when it has been selected, and is the new default resource.
3. The **Keyword Search** dialog box is displayed. If there are no search results for an **All Resources** search or you entered a null string, you will see a help message. A search on a specific resource with no results displays a resources tab with no results listed.
4. When search results are displayed, click the resource name hyperlink to go to the **Resources** tab for the resource. When viewing the detail pane for the resource (for example, a list of volumes for a host) you can use the **Export to CSV** button (or **Actions** menu) to generate a resource report for selected detail pane items (for example, all or some of the volumes for a host), or you can select detail pane items to perform an operation (for example, selecting a volume and editing the volume label).
5. If the current search results do not meet your needs, or you want to run multiple searches without closing the **Keyword Search** dialog box, you can repeatedly enter a new search string, and select the resource to search. To perform a criteria search for storage resources, click the **More Searches** button to start the Search dialog box.

Related concepts

- About searching HCS resources on page 491

Related tasks

- Criteria search on page 493

Criteria search

To perform a criteria search, locate the Search box on the upper right side of the GUI. The search box is always visible.

Note: A Tiered Storage Manager license is required to save searches, as described below.

Procedure

1. To start the **Search** dialog box, click **down arrow**, and select **More Searches**.
   - Search criteria are specified in the left pane.
   - Search results are displayed in the right pane.
2. Use **Search for:** to select one of the following storage resources:
   - Volumes
   - Parity Groups
   - DP Pools
   - Free Space
3. To the right of **Search for:**, select **All** storage systems, or select a specific storage system.
4. The **Basic** and **Advanced** tabs provide further refinements for your search. The selected storage resource (volumes, parity groups, etc.) affects the search criteria available under the **Basic** and **Advanced** tabs. This can be seen by changing the selected storage resource and reviewing available options.
5. To run your search, click **Search**. When results are displayed, display controls are provided for easier viewing, including tiling options, a maximize control, and the **Column Settings** button for controlling what data is displayed. Depending on the resource being displayed, you will see buttons for performing an action on selected resources. For example, when looking at a list of volumes, you will be able to **Allocate Volumes** (to a host), **Create Pool** (only for storage systems supporting this feature), **Edit Label** (volume labels, which can be searched on later), **Export to CSV** (reporting) or **More Actions** (more volume related tasks).
6. To save a search, click **Save**, and enter a name and description, then select **Save as a public search** (check box) to share the search with others. Note that saved searches are displayed according to the storage resource. For example, if you run a parity group search and save it, it only displays when the selected storage resource is **Parity Group**. Similarly, changing the storage resource to **Volumes** will display only saved volume searches.
7. To run a saved search, select the storage resource, and select the saved search.
8. To edit a saved search, run it, then click **Save** to start the edit dialog box. Leave **Name:** as is. You may edit **Description:** and change **Save**

as a public search, then click **Save** to close the edit dialog box. If you have a saved search selected and change the search criteria, run the search, click **Save**, edit the description etc., and click **Save**.

9. To delete a saved search, run the saved search. When the search completes, click **Delete Saved Search**.

Related concepts

- About searching HCS resources on page 491

Related tasks

- Keyword search on page 492

Related references

- Criteria search dialog box on page 494

Criteria search dialog box

Use this dialog box to refine the search criteria you use for the volumes, parity groups, DP pools, and free space.

After you have selected the storage resources and storage systems to search, you can choose the Basic tab or the Advanced tab options in this dialog box to refine your search.

Options differ with the selected storage resource as indicated by Vo (volumes), PG (parity groups), TP (DP pools), and FS (free space) in the following table.

When you enter the minimum required information in this dialog box, the Search button activates the search using the defined criteria. Click the Save button to name and save the search criteria you defined for later use.

The following table describes the dialog box fields, subfields, and field groups. In the Basic tab, a *field group* is a collection of fields that are related to a specific action or configuration. You can minimize and expand field groups by clicking the double-arrow symbol (>>).

The Advanced tab contains the following attributes:

- All available categories for a search type (volumes, parity group, etc.) can be combined into a compound search using the +/- buttons.
- You can create a compound search where All/Any conditions need to be met (and/or logic).
- A logical operator drop down (unlabeled) is provided for each category in your search. Values are text or symbols, and can include: in, not in, is between, is, is not, starts with, contains, =, >, >=, <, <=, and<>.
- The Value field provides either a drop-down list of values, or prompts you to enter data (names or numbers). When entering data, a valid entry displays a blue border, an invalid entry (format, spacing) displays a red border.

- Saved compound searches are re-displayed in total (all set categories, attributes and values) when you run your saved search. You can review, modify, and run new search variations, and save them.
- Use Column Settings to correlate displayed search results with your search. This helps validate that your search is returning the expected results.

Table 22 Search dialog box

Basic Tab	Description	Vo	PG	TP	FS
Volume Status	Defaults to Any, or you can select Allocated (in-use) volumes only or Unallocated (ready-for-use) volumes only.	x	-	-	-
Volume Type	Defaults to Any, or you can select Basic volume (a volume not from a DP Pool), or DP volume. Note that DP volumes use up pool space on demand, while basic volumes are created with a static space allocation.	x	-	-	-
>> Host Group	You can find volumes for any host, or filter by a selected host. For example, using any volume status and any volume type for a specific host finds all volumes for the host. You can refine your search with volume status (allocated, unallocated) and volume type (basic, DP).	x	x	-	-
>> Performance Criteria	This is Drive type (Any, FC, SATA, SAS, etc.), Drive speed (Any, 7200, 10000, 15000, etc.), and Chip Type (Any, SLC, MLC).	x	x	-	x
>> Availability Criteria	RAID Level (RAID 1, 5, 6, etc.) and Replication Function (Any, Simplex volume, Copy volume).	x	x	-	x
>> Capacity Criteria	Volume capacity (minimum or maximum in TB, GB, MB) and Used capacity (min, max in TB, GB, MB). For DP Pools, the field is Free capacity greater than, or less than a user supplied value in TB, GB, MB.	x	-	x	x
>> Subscription Criteria	For DP Pools only, Subscription, Subscription warning, Subscription limit are fields where a threshold is greater than or less than a user supplied percentage value. For example, find DP volumes subscribed at greater than 50%.	-	-	x	-
Advanced Tab	Description Note that the examples below describe all available categories, using volume search examples, except the final category search which applies to DP Pools only.	Vo	PG	TP	FS
Storage System	For a Volume search, using Attribute, you can specify Storage System Vendor, or Storage System Model (useful if you have multiple storage systems of the same type). The Value drop down provides available vendors or models. You can invert the search results using the is/is not operator. The +/- buttons allow the addition or removal of additional search conditions. When switching to other supported search types, investigate the Attribute and Value fields to see what is available.	x	x	x	x
ID/Number/ Name	For a Volume search, using Attribute you can specify a Parity Group, Pool ID, or LDEV ID as the means of locating volumes. The Value field prompts for ID information. The is between operator lets you find volumes in a range, for example, LDEV IDs 00:00:00-00:00:05.	x	-	x	x
Allocation	For a Volume search, using Attribute you can specify Allocated, then qualify it using a Yes/No selection in the Value field.	x	-	-	-

Volume Type	For a Volume search, using Attribute you can specify DP Volume Type, then qualify it with Basic, DP, or DP Pool in the Value field.	x	-	-	-
Label	For a Volume search, using Attribute, you can specify Label, then qualify it by entering a label value in the Value field. To use a partial text string, use the starts with or contains operators.	x	-	-	-
Performance	For a Volume search, using Attribute you can specify Drive Type, Drive Speed (RPM), or Chip Type. Drive types, drive speeds, and chip types are displayed in Value.	x	x	-	x
Availability	For a Volume search, using Attribute you can specify RAID Level or Copy type and role. RAID levels or replication types will be displayed in Value. Note there may be 20 or more replication types listed so you can be very specific in your search.	x	x	-	x
Capacity	For a Volume search, using Attribute you can specify Volume Capacity, Physical drive size, Used Capacity, or Used %. Enter a GB, TB, or MB value in Value.	x	-	x	x
Consumer	For a Volume search, using Attribute you can specify I/O Consumer Host, Storage Port, or Host Group/iSCSI Target. Enter a text value in Value. This category is meaningless if searching for unallocated volumes.	x	-	-	-
Security	For a Volume search, using Attribute you can specify Encryption, and then qualify it using a Yes/No selection in the Value field.	x	-	-	-
Migration	For a Volume search, using Attribute you can specify Volume Lock Status, and then qualify it using a Locked/Unlocked selection in the Value field.	x	-	-	-
Emulation	For a Volume search, using Attribute you can specify Emulation, and then qualify it by your selection in the Value field (for example, Open-V emulation).	x	-	-	-
Logical Partitions	For a Volume search, using Attribute you can specify CLPR or SLPR, and then enter a number, for example 1, in the Value field.	x	-	-	-
Subscription	For a DP Pool search, using Attribute you can specify Subscription (used space), Subscription Warning (warning threshold for volume creation), or Subscription Limit (threshold where volume creation is not allowed), and enter a number representing a percentage (you do not need to enter the % symbol). For example, you can find DP Pools with a subscription limit equal to 120%, or warning thresholds equal to or greater than 70%. Use Column Settings to ensure these attributes are displayed in your search results.	-	-	x	-

Related tasks

- Criteria search on page 493

Generating resource reports

This module describes how to export information about hosts, file servers, logical groups (public and private), user groups, and search results to a CSV file.

About generating reports

Information about registered storage systems and hosts and other items can be output to a CSV file. The CSV file can be used as a report, or used as a source file to generate customized reports.

The following items are examples of types of information that can be output to a CSV file:

- Host information
  - A list of hosts managed by Device Manager
  - Information about hosts for each OS
  - Information about volumes assigned to hosts
- File server information
  - A list of file servers managed by Device Manager
  - Information about file systems or file shares created on the file server
  - Information about volumes assigned to the file server
- Logical group information
  - A list of logical groups managed by Device Manager
  - Information about hosts or volumes in the selected logical group
- Resource search results
  - Information about volumes, parity groups, DP pools, and free capacity displayed in search results.
- Access control information
  - Relationship between user groups and users
  - Relationship between user groups, resource groups, and roles

Related tasks

- Exporting host information to a CSV file on page 497
- Exporting file server information to a CSV file on page 498
- Exporting logical group information to a CSV file on page 499
- Exporting search results to a CSV file on page 499
- Exporting access control information for resources to a CSV file on page 500

Exporting host information to a CSV file

Host information and a summary of the volume capacity allocated to hosts can be output to a CSV file.

Procedure

1. From the **Resources** tab, select **Hosts**.
2. Expand the tree and select the resource whose information you want to export to a CSV file.

- To export a list of all hosts managed by Device Manager, select **All Hosts**, and from the **Actions** menu located in the corner of the application pane, select **Export to CSV**.
- To export a list of hosts for a specific OS, select the target OS in the tree view, and click the **Export to CSV** button.
- To export host volume information, from the **Hosts** list click the host name link to list volumes, and click **Export to CSV**.

3. Click the **Save** button and specify the storage location for the CSV file.

Result

A CSV file is saved to the specified location.

Related concepts

- About generating reports on page 497

Exporting file server information to a CSV file

File server information, including assigned volumes and created file systems, can be exported to a CSV file. For cluster configurations, information is output for each cluster.

Procedure

1. From the **Resources** tab, select **File Servers**.
2. Select the resources whose information you want to export to a CSV file.
   - To output a list of all clusters and all file servers managed by Device Manager, select **All File Servers**.
   - If the file server belongs to the NAS Platform family (or NAS module), to output information about system drives assigned to the file server or NAS module, or a list of file systems on the file server or VSP G400, G600, G800, and VSP Fx00 models with built-in NAS modules, expand the tree, select the target cluster or file server or NAS module, and then select tabs according to the information that you need. To output a list of file systems or information about file shares for each Enterprise Virtual Server (EVS), expand the tree, select the target EVS, and then select tabs according to the information that you need.
   - If the file server is a NAS Platform F or Data Ingestor, to output information about volumes assigned to the file server or a list of file systems on the file server, expand the tree, select the target cluster, and then select tabs according to the information that you need. To output a list of file systems or information about file shares for each file server (node), expand the tree, select the target file server (node), and then select tabs according to the information that you need.
3. Click **Export to CSV**.
4. Click the **Save** button and specify the storage location for the CSV file.

Result

A CSV file is saved to the specified location.

Related concepts

- About generating reports on page 497

Exporting logical group information to a CSV file

You can save information about either of the supported logical group types (Public Logical Groups or Private Logical Groups) to a CSV file, including allocated volume capacity and assigned hosts and volumes.

Procedure

1. From the **Resources** tab, select **Logical Groups**.
2. Select the logical group whose information you want to export to a CSV file.
   - To output a list of all logical groups managed by Device Manager, select **Logical Groups**.
   - To output information about hosts or volumes included in logical groups, expand the tree, and then select the logical groups.
3. Select **Export to CSV** from the **Actions** menu located in the corner of the application pane.
4. Click the **Save** button and specify the storage location for the CSV file.

Result

A CSV file is saved to the specified location.

Related concepts

- About generating reports on page 497

Exporting search results to a CSV file

You can save your storage resource search results to a CSV file.

Procedure

1. From the **Search** pull-down menu, select **More Searches**, then specify search conditions and perform your search.
2. Click the **Export to CSV** button.
3. Click the **Save** button and specify the storage location for the CSV file.

Result

A CSV file is saved to the specified location.

Related concepts

- About generating reports on page 497

Exporting access control information for resources to a CSV file

You can export user-related information, such as access control information, to a CSV file.

Procedure

1. From the **Administration** tab, select **User Groups**.
2. From the **User Groups** tab or **Users** tab, click the **Export to CSV** button.
3. Click the **Save** button and specify the storage location for the CSV file.

Result

A CSV file is saved to the specified location.

Related concepts

- About generating reports on page 497

Using reports to verify system changes

You can use Device Manager - Storage Navigator to create, view, and delete reports. Reports contain information about your storage system's physical configurations and logical settings.

Reports can also provide a copy of the settings, which are useful when you intend to make changes. If you make a similar report after the changes, you can compare the reports to verify that new settings were made as intended.

Device Manager - Storage Navigator reports can be saved in a CSV or HTML file. Tables in the HTML version of the configuration report are sortable.

Viewing a Device Manager - Storage Navigator report

Before you begin

- Adobe Flash Player must be installed.
- Users can view the reports that they created.
- The window to specify the saving folder location may not be displayed when downloading the report in Google Chrome. In this case, follow Chrome Menu > Settings > Show advanced settings and uncheck the Protect you and your device from dangerous sites checkbox under Privacy.

Procedure

1. On the **Resources** tab, click **Storage Systems**, and then expand **All Storage Systems** and the target storage system.
2. Choose one of the following options.

- For VSP G series or VSP F series:
  Right-click the storage system, and then select **Reports**.
- For other available storage systems:
  From the **Actions** list in the application pane, select **Element Manager**. Refer to the documentation for the native management tool for your storage system.

3. Specify the report to download.
4. Click **Download Reports**.
5. Specify a folder in which to save a `.tgz` file.
6. Extract the downloaded `.tgz` file.
7. Display the report.
   For HTML reports:

   Open the file *extracted-folder*`\html\index.html`.

   For CSV reports:

   Open a CSV file in the folder *extracted-folder*`\csv`.

Viewing a report in the Reports window

You can view only HTML format reports in the **Reports** window. You can view CSV format reports in the previous procedure.

Procedure

1. On the **Resources** tab, click **Storage Systems**, and then expand **All Storage Systems** and the target storage system.
2. Choose one of the following options.
   - For VSP G series or VSP F series:
     Right-click the storage system, and then select **Reports**.
   - For other available storage systems:
     From the **Actions** list in the application pane, select **Element Manager**. Refer to the documentation for the native management tool for your storage system.
3. Click the name of the report to display.
   The report is displayed in the **Reports** window.
4. In the **Reports** window, click the name of the report in the list at the left, and then view the report at the right.

Creating a configuration report

You can create up to 20 configuration reports for each storage system. If you already created 20 reports, delete unnecessary reports first, and then create a new report.

Procedure

1. On the **Resources** tab, click **Storage Systems**, and then expand **All Storage Systems** and the target storage system.
2. Choose one of the following options.
   - For VSP G series or VSP F series:
     Right-click the storage system, and then select **Reports**.
   - For other available storage systems:
     From the **Actions** list in the application pane, select **Element Manager**. Refer to the documentation for the native management tool for your storage system.
3. Open the **Create Configuration Report** window.

   In Hitachi Command Suite:

   a. On the **Resources** tab, expand the **Storage Systems** tree.
   b. Right-click the local storage system, and then click **Reports > Create Configuration Report**.

   In Device Manager - Storage Navigator:
   - From **General Tasks**, click **Create Configuration Report**.
4. Specify a task name and click **Apply**. This task name is used as the report name in the **Reports** window. This process takes approximately 10 minutes to complete.
5. Click **Refresh** to update the **Reports** window. The created report appears in the list.

Deleting a configuration report

You can delete a report when you no longer need it, or to make room in the **Reports** window when the number of reports is near the limit.

Procedure

1. On the **Resources** tab, click **Storage Systems**, and then expand **All Storage Systems** and the target storage system.
2. Choose one of the following options.
   - For VSP G series or VSP F series:
     Right-click the storage system, and then select **Reports**.
   - For other available storage systems:
     From the **Actions** list in the application pane, select **Element Manager**. Refer to the documentation for the native management tool for your storage system.
3. Select the report to delete.
4. Click **Delete Reports**.
5. Click **Apply**.

14

Managing tasks

This module describes how to manage tasks.

- About tasks
- Customizing which tasks display in the global monitoring bar
- Viewing HCS task status
- Rescheduling HCS tasks waiting to be executed
- Stopping running data migration or data placement profile tasks
- Canceling scheduled HCS tasks
- Moving HCS tasks to the HCS Task History tab
- Restarting a failed or stopped HCS task
- Deleting HCS tasks
- Viewing system task status
- Managing system tasks
- Troubleshooting system tasks
- Viewing data collection task status
- Restarting a data collection task

About tasks

Tasks are submitted to perform specific functions, such as to allocate or create volumes, or to group storage resources. For tasks that you or other users submit, task status and task results display in a task list.

The method for submitting tasks and the tasks you can run differ depending on the task type. There are three types of tasks in HCS:

- HCS Tasks
  HCS tasks are tasks that are submitted using Hitachi Command Suite. Typically, you schedule HCS tasks to run at a time that is best suited for your storage system environment, such as during a period of low storage system activity. Resources connected to a task, such as volumes and parity groups, are reserved so that other tasks that use these resources do not run at the same time. The status of each registered task changes when the task finishes by completing successfully, failing, or stopping. Until a registered task finishes, you should not perform any other tasks on any volumes included in the task.
  From the Tasks & Alerts tab, HCS Tasks list, depending on the status of the task, you can:
  - Cancel tasks that are in Waiting status
  - Stop migration tasks or data placement profile tasks
  - Move finished tasks to the HCS Task History tab
    You cannot restart deleted tasks or tasks that have been moved to the HCS Task History tab. Additionally, HCS tasks that have been moved to the history are automatically deleted in order beginning from the oldest when the number of tasks exceeds 100,000.
  - Reschedule tasks that are in Waiting status
  - Delete completed tasks
  - Restart failed or stopped tasks
    To restart a task, it must meet the requirements that allow it to be restarted. Restarted tasks inherit the same configured values that were initially set when the task was registered and submitted. The configuration of the resource might have changed between the time when the task was initially submitted and when it is restarted, so when restarting a task, verify the configuration to determine whether you can restart it.
- System Tasks
  System tasks are specific to VSP Fx00 models, VSP Gx00 models or VSP G1000, VSP G1500, or VSP F1500 storage systems, and are submitted from Hitachi Device Manager - Storage Navigator.
  From the Tasks & Alerts tab, System Tasks list, you can view:
  - Task name
  - Task type
  - User

- Storage system
- Status (link)

  Clicking the status link will display additional status information in a **Status Details** window. For example, when a system task runs (Delete LDEVs for example) an Update Database sub-task that applies storage configuration changes to the HCS database also runs. When you click the status link of the primary task (Delete LDEVs), you will see the Update Database sub-task displayed in the **Status Details** window. Note that the Update Database task is also displayed in the Data Collection Tasks tab like other HCS data collection activities.

  To take further action on a system task, you would click Manage System Tasks and in the displayed window actions would include:
  - View task details
  - Delete or modify task settings
  - Apply configuration changes as a result of the task completion
  - Suspend or resume tasks (system tasks cannot be restarted)
- Data Collection Tasks

  Data collection tasks are listed on the Administration tab, under Managed Resources, on the Data Collection Tasks tab. These tasks can be sorted and filtered like tasks listed on the Tasks & Alerts tab.

  Data collection tasks are created by adding, refreshing, editing, and removing storage systems, file servers, and hosts from the Administration tab or the Resources tab.

  Data collection tasks are also created by submitting System tasks that perform an Update Database sub task.

  On the Data Collection Tasks tab, in the data collection tasks list, you can:
  - Check task results and task status
  - Delete tasks
  - Restart tasks

Note: Adding a file server is not registered as a data collection task because adding a file server is performed by registering Device Manager information in the management software of the file server.

Host information can be synchronized with Compute Systems Manager hosts. This task is listed as a data collection task.

- For HCS tasks and data collection tasks that do not involve VVols (where VVols are not used), you can receive an email notification at the completion of the migration task by specifying an email address when submitting tasks. Email notification is sent when the task completes and no longer needs its progress monitored; for example when the task ends successfully, fails, or stops. Use email notification when you want to verify the results of a task, when it is necessary to contact a server administrator or application administrator about changes to storage system settings, or when you submit a task that could take a long time to complete.

The global monitoring bar provides the following task functionality:

- Displayed at the bottom of the HCS GUI, these task status links provide convenient access for reviewing and managing HCS task and System task information from the Tasks & Alerts tab. The numbers displayed on the global monitoring bar are the sum of HCS tasks and system tasks. Note that data collection tasks cannot be accessed from the global monitoring bar. Task status links provide easy access to the following task categories:
  - Waiting (Waiting, Suspended)
  - In Progress (In Progress, Stopping, Stopping Immediate, Canceling, or Deleting, or Updating Database)
  - Completed (Completed, Canceled, or Stopped)
  - Failed (Failed or Database Update Failed)
- Clicking a task status link will display a dialog box where you can review and manage HCS tasks and System tasks on separate tabs. Filtering and display controls in the dialog box are the same as in the Tasks & Alerts tab. The management options available in the dialog box depend on the task status category. For example, you can stop an in progress HCS task, but the stop option would not be available for completed or failed HCS tasks.
- Depending on the category of tasks you select, you can also control which tasks are visible using a cut off date, or by moving tasks to task history.

Take the following actions to troubleshoot task errors:

- If, during the execution of a task, an error message displays indicating a database blockage, terminate the Device Manager server, and then recover the database. After recovering the database, restart the Device Manager server, and then refresh the storage systems. For the task that ended in an error, review the storage system logs and then, if necessary, run the task again.
- If a window does not display or if information to be displayed cannot be obtained, follow the instructions in the displayed error message, check the network status, and then run the task again or contact the server administrator.

Related tasks

- Viewing HCS task status on page 507
- Rescheduling HCS tasks waiting to be executed on page 508
- Stopping running data migration or data placement profile tasks on page 508
- Canceling scheduled HCS tasks on page 509
- Moving HCS tasks to the HCS Task History tab on page 509
- Restarting a failed or stopped HCS task on page 510
- Managing system tasks on page 512
- Viewing data collection task status on page 514
- Customizing which tasks display in the global monitoring bar on page 507

Customizing which tasks display in the global monitoring bar

You can control which tasks are visible when using the task status links by filtering tasks with a cut off date or moving failed tasks to task history.

The cut off date serves as a filter for the tasks from the Tasks & Alerts tab. You can change the date to restore older tasks to your view.

Moved tasks are can be displayed from Tasks & Alerts, HCS Task History tab.

Note: The Move to History button is also available in the Tasks & Alerts tab and applies to all task types, not just failed tasks. You may also select failed tasks and click Delete Tasks to permanently remove them.

Procedure

Specify a cut off date to reduce the number of displayed tasks, and/or move failed tasks to the **HCS Task History** tab, as follows:

a. Click **Failed** or **Completed** to display tasks. Click **Edit Duration** and specify a cutoff date, and click **OK**.
b. Click **Failed**, select HCS failed tasks, and click **Move to History**.

Result

The size of the visible task list is reduced.

Related concepts

- About tasks on page 504

Viewing HCS task status

View the status of HCS tasks to see if the task has completed, summary information about the task, and task configuration details.

Procedure

1. On the **Tasks & Alerts** tab, select **All Tasks**.
2. View the **Status** column on the **HCS Tasks** tab or on the **HCS Task History** tab to determine if a task has completed.
3. Click the task name link in the **Task** column on either the **HCS Tasks** tab or the **HCS Task History** tab to view the **Task Details** window, which provides the task details and task summary.

   If a task has failed, the **Message** field provides a description for why the task failed.

Related concepts

- About tasks on page 504

Rescheduling HCS tasks waiting to be executed

You may reschedule HCS tasks waiting to be executed.

Procedure

1. On the **Tasks & Alerts** tab, select **All Tasks**.
2. Select the waiting task you want to reschedule from the **HCS Tasks** tab and click the **Reschedule Tasks** button to specify a new schedule from the dialog box.
3. You can now either execute the rescheduled task or confirm that the waiting task is rescheduled by viewing the **HCS Tasks** list.

Related concepts

- About tasks on page 504

Related tasks

- Viewing HCS task status on page 507

Stopping running data migration or data placement profile tasks

Data migration or data placement profile tasks can be stopped during execution, in the following ways:

- Normal stop:
  A task stop request is issued and processing that is currently being performed is continued. For data migration, tasks are stopped when data moving processing completes. For Data Placement Profile operations, tasks are stopped when processing for each storage system completes. Processing that has not been started is stopped. It might take a long time after a task stop request is issued before the task stops completely. Stopped tasks can be restarted.
- Immediate stop:
  If you select the Stop the tasks immediately check box (immediate stop), or if a task is performed that searches for inconsistent data placement profiles when a storage system is refreshed, all processes are stopped immediately. For data migration tasks, since data that was being migrated remains in the target volume, the user must delete the data in the target volume. Therefore, we recommend that you use immediate stop only for emergencies.

The storage administrator checks the task progress and results in the Tasks & Alerts tab allowing the schedule of a task that is waiting to be changed.

Procedure

1. In the **Tasks & Alerts** tab, select **All Tasks**.
2. From the **HCS Tasks** tab, select one or more running tasks that you want to stop, and click **Stop Tasks**.
3. In the confirmation dialog box, click **OK**.
4. In the HCS tasks list window, confirm that the status of the task is stopped.

Note: To access the **Task Details** dialog box to see the point at which a task was stopped, click **All Tasks** in the **Tasks & Alerts** tree and click the task name link in the summary column of the **HCS Tasks** tab.

Related concepts

- About tasks on page 504

Canceling scheduled HCS tasks

You can cancel scheduled HCS tasks that have not yet started.

Note: You can only cancel tasks that are in Waiting status.

Procedure

1. On the **Tasks & Alerts** tab, select **All Tasks** and select a task from the HCS tasks list.
2. Click **Cancel Tasks**.

Result

Upon cancellation of a task the status of the task appears as Canceled.

Related concepts

- About tasks on page 504

Moving HCS tasks to the HCS Task History tab

You can move completed, failed, stopped, or canceled tasks from the HCS Tasks tab to the HCS Task History tab.

For example, you could manage the completed and cancelled tasks by moving them to the HCS Task History tab and keep the submitted, stopped, and failed tasks in the HCS Tasks tab until these tasks run to completion or are resolved.

After a task has been moved to the HCS Task History tab, it cannot be moved back to the HCS Tasks tab.

Note: After a task has been moved to the HCS Task History tab, the task cannot be restarted again.

Procedure

1. On the **Tasks & Alerts** tab, select **All Tasks**.
2. On the **HCS Tasks** tab, select one or more tasks and click **Move to History**.
3. Verify the tasks to move and click **OK**.

Result

The selected tasks are moved to the HCS Task History tab.

Related concepts

- About tasks on page 504

Related tasks

- Deleting HCS tasks on page 511

Restarting a failed or stopped HCS task

Any HCS tasks that have failed or stopped prior to completion can be restarted. Restarted tasks inherit the values that were set at the time the task was submitted.

Note: After a HCS task has been moved to the HCS Task History tab or has been deleted, it cannot be restarted. If a HCS task has been restarted several times, and the last instance of the task that you attempted to restart is deleted, the task cannot be restarted again.

Before you begin

Verify that the status of the target HCS task is Failed or Stopped.

Procedure

1. On the **Tasks & Alerts** tab, click the **HCS Tasks** tab, and select a **Failed** or **Stopped** task.

2. For the selected **Failed** or **Stopped** task, view the **Task Details** dialog box, and verify in the Message field whether the task can be restarted.
3. Click **Restart Task**.
4. Click **Submit** to restart the process.
5. You can check the progress and the result of the task on the **Tasks & Alerts** tab.
   Verify the results for each restarted task by viewing the details of the task.

Result

The status of the task you restarted no longer appears as Failed or Stopped in the list of HCS tasks.

Related concepts

- About tasks on page 504

Related tasks

- Viewing HCS task status on page 507

Deleting HCS tasks

You can delete tasks from the HCS Tasks tab or the HCS Task History tab.

Procedure

1. On the **Tasks & Alerts** tab, select **All Tasks**.
2. On the **HCS Tasks** tab or **HCS Task History** tab, select one or more tasks and click **Delete Tasks**.
3. Click **OK**.

Result

The selected tasks are deleted.

Viewing system task status

Viewing the status of system tasks allows you to see if the task has completed, as well as other details.

Procedure

1. On the **Tasks & Alerts** tab, select **All Tasks**.
2. On the **System Tasks** tab, view the **Status** column to determine if a task has completed.

3. To view task details, on the **System Tasks** tab, click **Manage System Tasks**, select the storage system for which the task was submitted, click **OK**, and review the task list.

Note: Click the **Status** column link to view the **Status Details** window, allowing you to view the progress of sub-tasks such as Update Database. Use the **Data Collection Tasks** tab to view details of the Update Database sub-task or to re-execute the task.

Related concepts

- About tasks on page 504

Related tasks

- Managing system tasks on page 512
- Troubleshooting system tasks on page 513

Managing system tasks

If your storage system is VSP Fx00 models, VSP Gx00 models or VSP G1000, VSP G1500, or VSP F1500, some tasks are registered as system tasks.

For system tasks, you can view task details, and suspend, resume, cancel, or prevent tasks from being automatically deleted.

Procedure

1. On the **Tasks & Alerts** tab, select **All Tasks**.
2. On the **System Tasks** tab, click **Manage System Tasks**.
3. Select the storage system where the task is registered and click **OK**.
4. Select a task and choose one of the following options:
   - Click **Delete Tasks** to delete or cancel a task.
   - Click **Suspend Tasks** to suspend a queued task.
   - Click **Resume Tasks** to resume a suspended task.
   - Click **More Actions > Enable Auto Delete** to delete completed tasks when the task list is full.

Note: HCS can manage many storage systems. Therefore you can manage up to 100,000 tasks before the oldest tasks start getting deleted. The system task list will display 384 tasks (for each storage system), including up to 128 tasks with a status of In Progress, Waiting, or Suspended, and up to 256 tasks with a status of Completed or Failed. Enable Auto Delete is recommended.

- Click **More Actions > Disable Auto Delete** to keep tasks after they are completed.

5. Verify the settings and click **Apply**.



Tip: The Tasks pane automatically updates every 60 seconds by default. To change the interval, from the **Storage Systems** tree, right-click the target storage system, and then select **Other Functions**. In the displayed window, click **Setting > Environmental Settings > Edit Information Display Setting**, then specify the refresh interval for the tasks list.

Related concepts

- About tasks on page 504

Related tasks

- Troubleshooting system tasks on page 513
- Viewing system task status on page 511

Troubleshooting system tasks

Verify the status of queued system tasks in VSP Fx00 models, VSP Gx00 models or VSP G1000, VSP G1500, or VSP F1500 storage systems that do not successfully complete within a reasonable amount of time and troubleshoot as necessary.

Check the details of a system task that has not completed successfully. System tasks cannot be restarted but system tasks can be suspended or resumed.



Note: This procedure is for troubleshooting problems for tasks related to changing a storage systems configuration, not for update database tasks, which are recorded in the Data Collection Tasks tab.

Procedure

1. On the **Tasks & Alerts** tab, select **All Tasks**.
2. Check the task **Status** column in the **System Tasks** tab to determine if a task has completed.
3. To view task details, click **Manage System Tasks**, select the storage system where the task is registered, and click **OK**.
4. Verify the task status and details by performing the following actions:
   - When a task status is **Failed**, click the link to view the reason. You can correct the problem and submit the task again. Click **Delete Tasks** to delete the task that failed.
   - When a task does not complete, it might be that the task list contained more than the maximum number of tasks allowed (384). To

resolve this problem, click **More Actions > Enable Auto Delete**. When **Enable Auto Delete** is set, tasks are automatically deleted after they successfully complete and exceed the maximum allowed number.

Note: HCS can manage many storage systems. Therefore you can manage up to 100,000 tasks before the oldest tasks start getting deleted. The system task list will display 384 tasks, including up to 128 tasks with a status of In Progress, Waiting, or Suspended, and up to 256 tasks with a status of Completed or Failed. Enable Auto Delete is recommended.

- When a task does not complete, it can be that the storage system configuration changed. To diagnose this, check whether the Resource Lock is displayed for a long time. If so, this can indicate that another program is still running. Wait until that program finishes and submit the task again.
- When the percentage of progress for an ongoing task is slow to indicate a change, it can be that another task operation is in progress. Continue to check on the status of the task until it completes.

Related concepts

- About tasks on page 504

Related tasks

- Managing system tasks on page 512
- Viewing system task status on page 511

Viewing data collection task status

Viewing the status of data collection tasks allows you to see if the task has completed, as well as other details.

Procedure

1. On the **Administration** tab, select **Managed Resources**.
2. On the **Data Collection Tasks** tab, view the **Status** column to determine if a task has completed.
3. To view task details, on the **Data Collection Tasks** tab, click the task **Status** link to review the **Data Collection Task Detail** dialog box. If a task has failed, look for a diagnostic message in the **Message** field.

Note: If you see Update Database tasks, these are the result of System tasks being submitted. For example, if a System task is

run to delete LDEVs, an Update Database sub-task to update the HCS database is also run and is recorded under **Data Collection Tasks**.

Related concepts

- About tasks on page 504

Related tasks

- Restarting a data collection task on page 515

Restarting a data collection task

Data collection tasks might fail due to a communication error or because the target storage system has an exclusive control lock in place. Data collection tasks that fail before they complete can be restarted. Restarted tasks inherit the values that were set when the task was registered.

Procedure

1. On the **Administration** tab, select **Managed Resources**.
2. On the **Data Collection Tasks** tab, review the **Status** column for failed tasks, and click the **Status** column link for the data collection task you want to restart.
   View the Message box in the **Data Collection Task Detail** dialog box to determine whether the task can be restarted.
3. Click **Restart Task** and verify the information in the summary list of the **Restart Task** dialog box.
4. Click **Submit** to restart the task.
5. Verify the task status on the **Data Collection Tasks** tab.
   Verify the results for each task by viewing the details of the task.

Result

The status of the restarted data collection task no longer appears as Failed in the list of tasks.

Related concepts

- About tasks on page 504

Related tasks

- Viewing data collection task status on page 514

15

Monitoring managed resources and resolving alerts

This module describes how to monitor managed resources and alerts in a Storage Area Network (SAN).

- ☐ Monitoring storage resources
- ☐ Viewing information summaries
- ☐ Analyzing storage system performance
- ☐ Analyzing Universal Replicator performance
- ☐ Managing alerts

Monitoring storage resources

You can check the status of storage resources to detect problems and output reports to a CSV or PDF file containing search results or configuration and capacity information about storage systems, hosts, volumes, logical groups, and other storage resources.

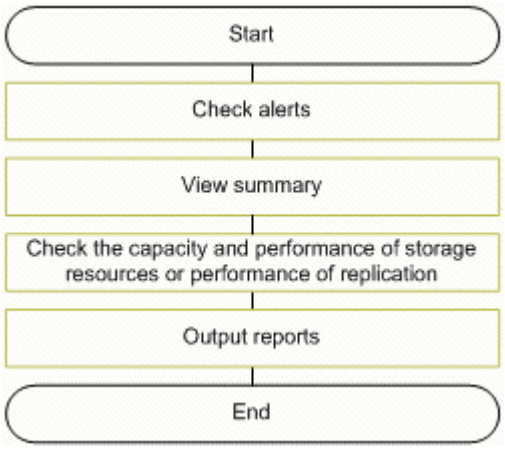
If you have the Replication Manager license, from the Replication tab, regularly check the status of the replication configuration and the performance information.

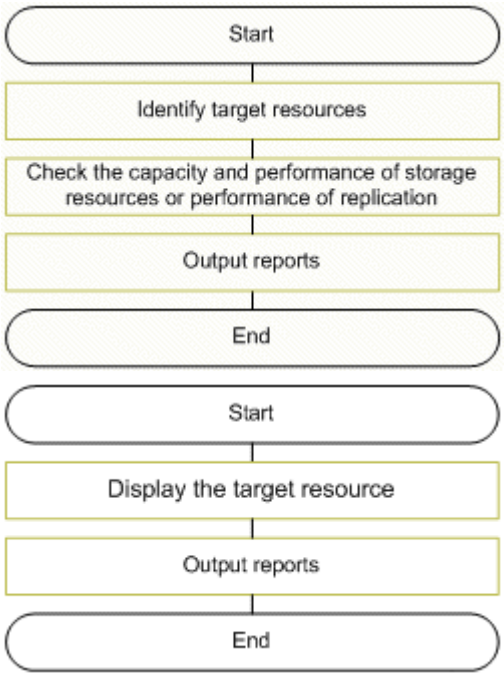
If you have the Tiered Storage Manager license, you can see the operational status of storage resources from the Mobility tab and perform integrated operations to optimize data placement.

If you have the Tuning Manager license, you can analyze storage system performance in the Analytics tab. You can identify whether the cause of a performance problem exists in storage systems in response to inquiries from the host and application administrators, and analyze whether storage resources are being used appropriately.

If you have a license for both Replication Manager and Tuning Manager, you can check the C/T delta performance of managed copy groups from the UR Performance Analysis window of the Replication tab. This allows you to analyze the cause of C/T delta degeneration and check how to resolve problems.

This information can be used to migrate volumes or add DP pool volumes to optimize storage resources. The following graphics illustrate the flow of monitoring and checking the status of storage resources. To check daily operational status, inquiries from a host or application administrator, see the following respective figures:





In summary, you can:

- View summary information from the Dashboard tab.
- Check alerts from the Tasks & Alerts tab or the dashboard.
- Search storage resources from the Search box on the tabs bar.
- Reference data migration information from the Mobility tab.
- Analyze performance from the Analytics tab.
- Analyze Universal Replicator performance from the UR Performance Analysis window of the Replication tab.

Viewing information summaries

This module describes how to view summary information for registered storage systems and hosts.

About checking system status in the dashboard

The Hitachi Command Suite dashboard shows a summary of tasks, alerts, and activities of registered storage systems.

- **Storage System Alerts:** A list of links for each registered storage system or file server and the associated alerts. You can choose particular severity levels of alerts to be collected in the Storage System Alerts report by clicking the tool icon in the application pane. Note that SNMP traps received from the file server are displayed as alerts. For details about how to specify settings to receive SNMP traps, see the *Hitachi Command Suite Administrator Guide*.
- **Datacenter Utilization - Physical capacities:** A bar chart that provides volume capacity information for Open-Allocated, Open-Unallocated, Open-

Reserved, Mainframe, and Free volumes. This chart applies to all of the managed storage systems, collectively. Because capacities are calculated based on actual volume capacities, the capacity of virtual volumes, such as DP volumes, is not included. If the drive type is FMC and if accelerated compression is enabled for the parity group, then the Physical Capacity Total will be added to Open-Reserved.

- **Failed HCS Tasks:** A list of failed HCS tasks submitted by you or other users. Each item includes the task name, task type, and description for the failed task. Clicking the Task column link for a task that you submitted allows you to view task details, including error messages. You may not view task details for tasks that you did not submit.
- **Failed System Tasks:** A list of failed System tasks submitted by you or other users. Each item includes the task name, task type, and description for the failed task. Clicking the Status column link for a task that you submitted allows you to view the processing status.
- **Top 10 Consumers:** A list of the ten hosts or file servers or NAS modules having the largest combination of allocated physical and virtual storage capacity. For file servers in a cluster configuration, if the file servers are being managed by System Management Unit v10.1.3070 or later, or by File Services Manager v3.2 or later, information is displayed by cluster.
  - The Capacity column reports (capacity of basic volumes + subscribed capacity of DP volumes).
  - The Used Capacity column reports (capacity of basic volumes + used capacity of DP volumes).
  - File System Used Capacity indicates the file system space used by file servers or NAS modules or hosts as (the total value calculated from File system Used % for each volume).

Note: The File System Used Capacity is displayed only for hosts that can collect information by using Host Data Collector or Device Manager agent , or for file servers.

This value cannot be collected for a virtualization server.

You can click host or file servers or NAS modules links on the Resources tab to see volume details related to capacity, used capacity, and volume type. By reviewing this capacity information, you can efficiently use storage resources and determine whether to migrate data to DP volumes.

- **Tier Residency:** A chart displaying the usage of hardware tiers for each logical group. You can compare the operating information of a logical group and the usage of hardware tiers by checking whether the hardware tiers for the logical group are configured appropriately. You can also select logical groups whose configuration is to be displayed. Click the tool icon in the application pane to configure logical groups that display the usage status of hardware tiers.

- The following file server or NAS module reporting is available in the Hitachi Command Suite dashboard, when the file servers or storage systems with NAS modules are registered in Device Manager:
  - **Top 10 File Systems:** When HCS is linked with a file server (NAS Platform family, NAS Platform F (HNAS F) or Data Ingestor (HDI)), or when or storage systems with NAS modules are registered, information about file systems that have a large capacity usage rate or free area can be checked to see if there is insufficient free space. By default, the allocated capacity usage rate is displayed in descending order. Clicking the column header of other items displays the top 10 file system information again, ordered by the selected item. You can also change the items displayed by clicking the tool icon in the application pane.
  - **Top 10 File Server Storage Pools:** When HCS is linked with a file server (NAS Platform family) or NAS module, information about storage pools that have a large capacity usage rate can be checked to see if there is insufficient free space. By default, the capacity usage rate is displayed in descending order. Clicking the column header of other items displays the information again, ordered by the selected item.
  - **Top 10 File System Snapshots:** When HCS is linked with a file server (NAS Platform F), information about volumes having a high capacity usage rate can be checked to see if there insufficient free space.
  - **File Level Tier Residency:** When HCS is linked with a file server (NAS Platform F), a graph is displayed for each file system tier, showing the usage status of the file systems and the volumes for file snapshots. You can review the usage rate and the used capacity and then expand the file system to increase free space and optimize the used capacity distributed among tiers. If you click the tools icon to start File Services Manager, you can review the policy settings on the tiered file system.

Related tasks

- Accessing the dashboard on page 521
- Customizing the dashboard on page 522

Accessing the dashboard

When you log in to HCS, the Dashboard tab and reports are displayed by default.

The Dashboard tab reports include:

- Datacenter Utilization - Physical Capacities
- Top 10 Consumers
- Top 10 File Systems
- Top 10 File Server Storage Pools
- Storage System Alerts
- Failed HCS Tasks
- Failed System Tasks
- Tier Residency

Click Dashboard Settings to select reports to display.

The title bar of each report has a right arrow icon [–>] that when clicked, takes you to another tab to review related detail information.

For example, clicking the arrow icon for the **Storage System Alerts** report displays the Tasks & Alerts tab with summary and detail information for alerts.

Click the Dashboard tab to re-display the dashboard reports.

When viewing Dashboard reports, click links to review specific detail.

For example, in the **Top 10 Consumers** report, clicking a specific host or file-server link displays the Resources tab, expands the hosts or file-server tree, and displays related summary and volume information.

Related concepts

- About checking system status in the dashboard on page 519

Related tasks

- Customizing the dashboard on page 522

Customizing the dashboard

Select dashboard reports, position reports, and change the column layout.

Click Dashboard Settings to:

- Select reports you want to display using the check boxes. Select all reports using the column header.
- Clear reports using the check boxes. Clear all reports using the column header.
- Choose the number of display columns for reports.
- Restore the default dashboard by setting No. of Dashboard Columns to 2 and checking Report Name to select all of the reports.
- HCS will retain changes across login sessions.

Using the title bar of each report:

- Re-locate a report by dragging the title bar to a new location.
- Double-click the title bar to maximize the report.
- Double-click the title bar again to minimize the report.
- You can also use the maximize/restore button.
- There is an up-arrow/down-arrow button that hides a report except for the title bar, providing more area for other reports.
- Use the close button to remove the report. To display it again, use Dashboard Settings.
- If shown, the tool button (a wrench) is for configuration of the report.

Related concepts

- About checking system status in the dashboard on page 519

About the Storage Systems tree

Hitachi Command Suite lists registered storage systems and provides information summaries in the Storage Systems tree on the Resources tab.

The Storage Systems tree lists:

- Storage systems
- DP Pools
- Parity groups
- Volumes
  - Open-Allocated
  - Open-Unallocated
  - Open-Reserved
  - Mainframe-Unspecified
- External Storage

The Storage Systems tree is the starting point for changing the configuration of these resources.

For VSP Gx00 models, VSP G1000, VSP G1500, VSP F1500 and VSP Fx00 models, additional functionality for configuring the storage system is available from Storage Systems tree. Use the right-click menu from the list of storage systems on the Resources tab, or click the System GUI link on the application pane to access the additional functionality. VSP Gx00 models and VSP Fx00 models also support launching the maintenance utility from the tree.

Note: For VSP Gx00 models, VSP G1000, VSP G1500, VSP F1500 and VSP Fx00 models, if you want to improve the security when Device Manager - Storage Navigator is started, set a value for the `client.launch.em.secure` property in the `client.properties` file. For details about this property, see the *Hitachi Command Suite Administrator Guide*.

Related tasks

- Viewing current storage system information on page 523

Viewing current storage system information

Configuration information for all managed resources is available on the Resources tab.

Procedure

1. On the **Resources** tab, select **Storage Systems**.

2. Select the parent group from the **Storage Systems** tree to view a volume summary and additional detailed information in the application pane.
3. Click on the link for a storage system to access additional details.
4. Continue to view more details by clicking the links on subsequent views.

Related concepts

- About the Storage Systems tree on page 523

Viewing MP Blade information

You can view and edit the MP Blade information.

Procedure

1. On the **Resources** tab, expand the **Storage Systems** tree, and select a target storage system.
2. Choose one of the following options:
   - For VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, and VSP F1500 storage systems:
     Select **Components**.
   - For other available storage systems:
     From the **Actions** list in the application pane, select **Element Manager**, and then from the **Storage Systems** tree, select **Components**.
3. On the **Components** window, click **Chassis ID**.

Result

The MP Blade information displays.



Tip: For details about editing MP Blade information, click Help.

About the Hosts tree

Hitachi Command Suite lists registered hosts by operating system in the Hosts tree on the Resources tab.

You can register both physical hosts and virtualization servers. You can also group the WWNs or iSCSI names that are used by specific storage systems, and set up that group as a host. You select a host from the Hosts tree to allocate or unallocate volumes. Note that you can reference all registered hosts, but for the list of volumes that are allocated to each host, you can only reference volumes that fall within your resource group assignments.

Related tasks

- Viewing current host information on page 525

Viewing current host information

Information for all managed hosts is available on the Resources tab.

Procedure

1. On the **Resources** tab, select **Hosts** to expand the tree.
2. Select an operating system to see an information summary and a list of hosts on the application pane.
3. Click the link for a host to access additional details.
4. Continue to view more details by clicking the links on subsequent views.

Related concepts

- About the Hosts tree on page 524

About the file server view

You view information about the volumes and file systems of managed file servers through the file server view.

Information acquired by the Device Manager can be hierarchically displayed by registering a file server as a Device Manager resource. Users can select a file server from the file server view to allocate or unallocate.

You can execute setting operations such as allocation and unallocation by cluster or by file server (node).

Related tasks

- Monitoring file servers on page 525

Monitoring file servers

You can check on the status of your file servers and review volume and file system information.

Procedure

1. On the **Resources** tab, select **File Servers**.
2. Expand the tree and select the resource whose information you want to browse.
3. In the application area, select the resource you want to reference.
4. Browse the information such as volumes and file systems.

Related concepts

- About the file server view on page 525

Analyzing storage system performance

This module describes procedures for analyzing storage system performance to determine if a storage system is the source of application performance problems.

About analyzing storage system performance

The Analytics tab provides information for analyzing storage system performance issues, and assists storage administrators in monitoring short term and long term trends.



Note: A Tuning Manager license is required to use the Analytics tab. For details on Tuning Manager, see the *Hitachi Command Suite Tuning Manager User Guide*.

If you are using the VMware VVol management functionality, note the following:

- When analyzing storage performance on a host basis, information about SLUs is not included. For this reason, you might not be able to identify the hosts that are increasing the load on the storage system.
- SLUs are not included in logical groups. Within the health check report, the "Affected Logical Groups" only display logical groups that include volumes that are related to resources that exceed the threshold value, and are not SLUs.

The following notes apply if the storage system (except the VSP G1000 (microcode version is 80-05-01-XX/XX or later), VSP G1500, and VSP F1500) to be analyzed contains a flash drive (FMD, FMC, and SSD):

- The particular characteristics of this drive type make it impossible to determine load status from the busy rate of a flash drive parity group and the usage rate of controllers. A large value might be displayed regardless of load. For this reason, a caution or warning status might be displayed in the following analysis results:
  - Parity groups page in the Identify Performance Problems wizard
  - The drives page in the Identify Performance Problems wizard
  - Busy PG by Utilization in Health check reports
  - PG Utilization (Threshold Exceeded) in Health check reports

Use the Analytics tab to check storage system performance:

- Identifying performance problems for storage systems:
  - If there is an application performance issue, obtain the information needed to identify a logical group. For example, this could be the host where the problem occurred, the label of the volume, or the mount point. Based on the acquired information, specify the period (up to the past 30 days) for analyzing the performance and the logical group, the

host, or the virtual machine that corresponds to the application in order to analyze the performance issue. Check that no resources exceed the threshold values and determine whether the cause of the performance issue is in the storage system.

- Analyzing and load balancing MP Blades or MP Units:
  - Identify MP Blade/Unit busy rates that exceed threshold values, and view resources placing a heavy load on the MP Blade/Unit. If possible, reduce the processing load or change the execution time of the processing using the resources. If you want to allocate resources to a different MP Blade/Unit, you can create a plan and simulate the effect of assigning resources to another MP Blade/Unit with a lower busy rate. To implement the plan, generate a raidcom command script based on the plan, and run the script using CCI to re-assign the resources.



Note: You can analyze MP Blades or MP Units only for VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, VSP F1500, Virtual Storage Platform, or HUS VM storage systems. The term MP Blade is used for VSP G1000, VSP G1500, VSP F1500, and Virtual Storage Platform. The term MP Unit is used for VSP Fx00 models, VSP Gx00 models and HUS VM.

- Analyzing storage performance on a host basis:
  When a performance problem occurs, for example when you receive notification from a host administrator or application owner, or when you receive health check system alerts, you might want to analyze the storage system performance on a host basis to determine if any hosts are increasing the load on the storage system. Identify the hosts with performance problems based on response time metrics, and plot a trend chart for response time metrics. Identify the hosts that are increasing the workload on the storage system based on IOPS and Data Transfer Rate metrics and plot trend charts for IOPS and Data Transfer Rate metrics. Compare the trend charts to determine which hosts are affecting the storage system performance.
- Performing periodic performance analysis of storage systems:
  - Analyze the performance of storage systems to periodically check that resources are being used appropriately.
  - The target storage system can be automatically analyzed at a specific interval; for example, weekly or monthly. For multiple storage systems, analysis schedules can be set to execute in a batch process to reduce the workload of storage system management. In addition, you can specify settings for performance analysis to send notifications by email of any errors that occur during analysis, as well as attach this information in the form of PDF-format analysis results reports to an email message.
  - By using these analysis results reports, you can check whether a resource exists that exceeds or nearly exceeds the threshold value. If

there is a problem in the resource, check the detailed performance information from the report that corresponds to each resource metric. Based on performance information such as the number of I/O operations (IOPS) and data traffic displayed in the report, discuss what measures to take and perform any volume allocation or data migration tasks. Resolving performance problems can prevent errors.

- By referring to past analysis results, problems can be specified and actions can be determined more easily. The analysis report can be output to a PDF or CSV file. If output to a PDF file, the contents and a summary are created so that the report can be submitted as is. If output to a CSV file, the report is useful as a general-purpose file, and it can be effectively used to create materials by editing the format.

For more detailed performance information, by starting Performance Reporter, you can perform an analysis in minute units or analysis of long-term trends. If you import a report definition to Performance Reporter in Tuning Manager in advance, you can display a report with items that correspond to the items in the Identify Performance Problems wizard, the Health Check report, or the Analyze MP Blades/MP Units window.

An appropriate default value is set for the threshold values to judge if there are problems with storage system performance. You can change the threshold values for their environment if necessary. In accordance with the specified threshold values, the background of the displayed analysis result chart is colored for each threshold range. The Caution threshold range is displayed in yellow. The Warning threshold range is displayed in pink.

Related tasks

- Identifying storage system performance problems on page 530
- Analyzing and load balancing MP Blades or MP Units on page 532
- Analyzing storage performance on a host basis on page 534

Notes on health check reports

If your site is using health check reports, you can refer to this section for additional information about Utilization Reports, health check report categories, and troubleshooting.

Review the following notes about health check reports:

- The table within a health check report shows the average of the values and the maximum values plotted in the chart. The values plotted in the chart are the average values across each unit specified for Interval, excluding the next report. In the following reports, the maximum value within the unit specified for Interval is shown:
  - PG Utilization (Threshold Exceeded)
  - CLPR Max Write Pending Rate by MPB/MPU

- The following categories are part of the health check reports, but are not displayed when using the Analytics tab. These health check report

categories represent storage system total front-end and back-end metrics that can be used with the Performance Reporter metrics to get a better understanding of storage system performance:
    - Storage system total front-end IOPS: represents the total number of IOPS (inclusive reads and writes) issued to a storage system by its connected hosts.
    - Storage system total front-end MB/s: represents the total number in megabytes per second that are read from or written to a storage system by its connected hosts.
    - Storage system total back-end IOPS: represents the part of the storage system total front-end IOPS that was not served by the cache memory in the storage system.
- A logical group that includes a volume used for a Thin Image copy pair with the clone attribute or a Thin Image copy pair that supports cascade configurations (a copy pair whose secondary volume is a DP volume) might be displayed in **Affected Logical Groups**. If the logical group is displayed, use Performance Reporter to check the details.
- If Partially Unknown is displayed for Health Status of the health check reports, check the conditions, and apply the applicable solution, as outlined in the following table.

Conditions	Solutions
All of the following conditions apply: • The Storage System that is a target for Health Check has a Mainframe Volume. • The "Mainframe Volume Monitoring" setting is "N" or "n" in the instance settings in Tuning Manager - Agent for RAID.	Confirm that the settings of "Mainframe Volume Monitoring" is "Y" or "y" in the instance settings in Tuning Manager - Agent for RAID.
All of the following conditions apply: • The Storage System that is a target for Health Check has an Unassigned Open Volume. • The "Unassigned Open Volume Monitoring" setting is "N" or "n" in the instance settings in Tuning Manager - Agent for RAID.	Confirm that the settings of "Unassigned Open Volume Monitoring" is "Y" or "y" in the instance settings in Tuning Manager - Agent for RAID.
A parity group that has no volumes, exists. To confirm the volumes that have been created from the parity group, from the Resources tab, in the list of parity groups for the relevant storage system, check No. of Volumes. Repeat this check for all parity groups for the relevant storage.	Create volumes from the parity group.
The version of Tuning Manager - Agent for RAID is earlier than 5.9.0.	Upgrade Tuning Manager - Agent for RAID to version 5.9.0 (the version that supports Health Check) or a later version.

Identifying storage system performance problems

With Tuning Manager installed, you can identify storage system performance problems from the Analytics tab. You can check if resources that exceed the threshold values exist on a logical group, host, or virtual machine basis.

Before you begin

- Install Tuning Manager and link it to HCS.
- Register the Tuning Manager license.
- Import the report definition file in Tuning Manager before analyzing performance with the Analytics tab (ensure that you delete the report definition file for the previous version before importing the latest report definition file).
- Obtain the most recent configuration information for the storage system (if performance is to be analyzed on a virtual machine basis and VVol management functionality is used).

  When a virtual machine (guest OS) restarts, a VVol for swap might be re-created. Restarting a virtual machine (guest OS), or adding or removing disks to and from a virtual machine might change configuration information for the storage system. In this case, if you do not update the storage system information in the HCS system, the resources relevant to the SLUs allocated to the virtual machine might not be correctly displayed. Update the storage system information as needed.
- Specify the settings for acquiring the correspondence between a virtual machine and SLU (VVol), if performance is to be analyzed on a virtual machine basis and VVol management functionality is used.
- Identify the date and time of the performance problem.
- Identify application information such as host or logical group name (if performance is to be analyzed on a logical group or a host basis).
- Identify names of a virtual machine and a virtualization server (if performance is to be analyzed on a virtual machine basis)

For information about linking with Tuning Manager and importing the report definition file in Tuning Manager or for specifying settings for acquiring the correspondence between virtual machines and SLUs (VVols), see the *Hitachi Command Suite Administrator Guide*.

Procedure

1. On the **Analytics** tab, expand **Logical Groups**, **Hosts**, or **Virtual Machines** tree. If a virtual machine (guest OS) is to be analyzed, select the target virtual machine from the **Virtual Machines** tree.
2. From the displayed lists, select the target logical group, host, or virtual machine, and then click **Identify Performance Problems**.

   You can also select some of the volumes or hosts in a logical group.
3. Read the **Introduction** and click **Next**.

4. Based on the application information acquired from the host administrator, set the storage system type, volumes, and period to analyze, and then click **Next**.

   Note: Do not include Copy-on-Write Snapshot pair volumes, Thin Image pair volumes, or V-VOLs in the targets to be analyzed. Otherwise, you might be unable to correctly acquire some performance information. For this reason, you cannot determine the load status of the storage system.

   In the **Target Selection** window, if more than one volume type is listed in the drop-down list box of Step 1, repeat steps 4 to 6 for each volume type.
5. In **Overview**, check for resources that have exceeded the performance threshold values and also the metrics of the resources related to the affected resources, and then click **Next**.

   Tip: If there is no resource that exceeds a threshold value, it is likely that the problem does not originate from the storage system. If this is the case, close the Wizard.

6. If there is a volume or pool whose response time is longer than the expected response time, click **Next** to identify any resources that affect the volume or pool. If there is no relevant resource, you can assume that the storage system has no problem.
7. Check the information for each metric displayed for each type of storage system type, and then click **Next**. Check the analysis report. If needed, you can export the analysis results to a PDF file. The analysis result chart displays the Warning threshold range in pink.
   - For an MP Blade/Unit with threshold cautions, click **Analyze MP Blades/Units** to analyze MP Blade/Unit busy rate metrics and create a plan for reassigning resources to a less busy MP Blade or MP Unit to load balance host I/O requests.
   - For other resources that exceed a threshold value, there might be a problem with the storage system. If this is the case, click **Next**.

   If you analyze enterprise-class storage (VSP G1000, VSP G1500, VSP F1500, Virtual Storage Platform, Universal Storage Platform V/VM), HUS VM, VSP Fx00 models, or VSP Gx00 models, and when all volumes that belong to a parity group are included in the following list of volumes, the chart displaying Read Hit rate is not displayed:
   - Thin Image pool volume
   - DP pool volume

   Tip: To view detailed performance information, start Performance Reporter. From Performance Reporter, display the report by selecting **Identify Performance Problems** under **Identify**

Performance Problems and Health Check. You can analyze on a per-minute basis or on a long-term trend basis by setting the appropriate display conditions.

For details on using Performance Reporter, see the *Hitachi Command Suite Tuning Manager User Guide*.

Related concepts

- About analyzing storage system performance on page 526
- About global-active device on page 334
- Working with VMware VVols on page 406

Related tasks

- Creating periodic health check reports for storage systems performance analysis on page 540
- Analyzing and load balancing MP Blades or MP Units on page 532
- Analyzing storage performance on a host basis on page 534
- Acquiring the most recent storage system information on page 64

Related references

- Configuring the Performance Reporter display on page 548

Analyzing and load balancing MP Blades or MP Units

On the Analytics tab you can identify MP Blades and MP Units with a busy rate exceeding caution thresholds, and assign resources to an MP Blade or MP Unit with a lower busy rate to improve host application performance or balance I/O requests for disks with varying performance levels (HDD, SSD, FMD, FMC).

Before you begin

- Install Tuning Manager and link it to HCS.
- Register the Tuning Manager license.
- Import the report definition file from Tuning Manager before analyzing performance with the Analytics tab (ensure that you delete the report definition file for the previous version before importing the latest report definition file).
- To analyze MP Blade or MP Unit busy rates, the storage system must be a VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, VSP F1500, Virtual Storage Platform, or HUS VM. For more information, see the *Hitachi Command Suite Administrator Guide*.
- The storage administrator must have permission to log in to the server where CCI is installed, and permission to use `raidcom` commands to modify the target storage system so that MP Blade or MP Unit resource assignments can be completed.

Procedure

1. On the **Analytics** tab, choose an MP blade to analyze by doing one of the following:
   - Select a target host or logical group and click **Identify Performance Problems**. From the Overview page (if MP blade cautions appear) or from MP Blade/Unit page, click **Analyze MP Blades/Units**.
   - Expand the storage system tree, select a storage system from the application pane, and then click **Analyze MP Blades/Units**. The **Analyze MP Blades/Units** window opens.
2. On the **Analyze MP Blades/Units** window, verify that the **Analysis Target** information is correct.
3. In **MP Blades/Units Busy Rate**, examine the busy rate graph and determine which MP Blade is a candidate (based on exceeding thresholds, or having the highest busy rate) for reassigning resources to a less busy MP Blade. If a target MP Blade is not already selected, select the target MP Blade and click **Choose as Target MP Blade/Unit**.
4. In **Selected MP Blades/Units**, identify the resources causing the high busy rate.

   Note: In the list of resources in **Selected MP Blades/Units**, pair volume information is displayed by copy pair. Note that, for Thin Image copy pairs with the clone attribute or Thin Image copy pairs that support cascade configurations (copy pairs whose secondary volume is a DP volume), information might be displayed by volume. Select all volumes, and then complete a simulation.

5. After identifying busy resources for the target MP Blade, return to **MP Blades/Units Busy Rate** and select the destination MP Blade to which you want to assign resources, and then click **Choose as Destination MP Blades/Units**.
6. Select resources and click **Move** to assign resources to the destination MP Blade and create the assignment plan.

   Tip: For target MP Blade resources that should be excluded from the assignment plan, select the resource and click the **Lock** button. For example, system disks or disks that cannot afford a temporary deterioration in response time due to cache operations incurred when reassigning the resource to another MP Blade.

7. Verify that the created plan is acceptable, and then click **Generate Script Commands** to create the `raidcom` command script.
8. Log in to the server where CCI is installed.
9. Execute the `raidcom` command script from CCI to physically reassign resources.

10. Run Performance Reporter to check that the busy rate for the target MP Blade and destination MP Blade are acceptable.

Tip: To view detailed performance information, start Tuning Manager from the **Tools** menu. By displaying Performance Reporter reports from Tuning Manager, you can analyze information on a per-minute basis or based on long-term trends. For details on how to use Performance Reporter, see the *Hitachi Command Suite Tuning Manager User Guide*.

Result

The busy rate for the target and destination MP Blades is adjusted by the reassignment of resources.

Related concepts

- About analyzing storage system performance on page 526

Related tasks

- Identifying storage system performance problems on page 530
- Changing performance threshold settings to correctly analyze storage system performance on page 546

Related references

- Configuring the Performance Reporter display on page 548

Analyzing storage performance on a host basis

On the Analytics tab, you can analyze storage system performance by identifying the hosts that are increasing the workload on the storage system. You can plot the trend charts for response time, IOPS, and Data Transfer Rate metrics, and then compare the trends to identify which hosts are increasing the workload.

Note:

- Hovering your mouse in the plot area displays a reference line (red vertical line) which lets you view the date and time for a data point. You can show or hide the reference line by right-clicking the mouse.
- The host performance metrics are calculated by aggregating the performance values of all the LDEVs connected to the host. If the target host is sharing and managing an LDEV with another host, the performance values of the target host might include the performance values of the host with which the LDEV is shared. For example, the value displayed for each file server represents the whole cluster because the file server share the volume.

Before you begin

- Install Tuning Manager and link it to HCS.
- Register the Tuning Manager license.
- Verify if the hosts are registered in HCS.

Procedure

1. In the HCS search box, enter the target host name to perform a keyword search.
2. From the returned search results, click the link for a host to display a list of volumes and information about the storage systems used by that host. If there is more than one storage system listed, you might want to analyze each storage system.
3. On the **Analytics tab**, select the storage system for which you want to analyze performance, and click **Storage Performance Analysis per Host**.
4. In the **Storage Performance Analysis per Host** window, identify the hosts with performance problems based on response time metrics.

   To create a chart for average response time metrics:

   a. In the **Hosts** table, sort the hosts by Average Response Time.

      The hosts and file servers (by node) and NAS modules (by cluster) are listed in the **Hosts** table. The table also includes the host groups that are being used to externally connect storage systems.

   b. Select the hosts with high response time values.

      You can use the Worst Response Time metrics to analyze the peak workload periods.

   c. Select **New Analysis View**, and click **Show Chart**.

      The data points are plotted for the default time period. The default time period is 3 to 4 hours.

   d. In the performance trend chart displayed, you can view how the response times vary. The trend analysis helps you determine the variations in the response times and if the response times are longer than desired.

   e. (Optional) Change the time period and repeat the analysis to see how the response time data changes over time and how it affects the overall performance of the host applications.

      To change the time range, select the start and end dates, select **Data Granularity**, specify the number of minutes, hours, or days, and then click **Apply**.

      You can also customize the time period settings by clicking the **Visual Settings** button, which lets you adjust the time range using the slider. To change the time range using the slider, one or more performance charts must be displayed in an analysis view.

Note: When analyzing the performance data in minutes, specify a time period no longer than 24 hours. When the specified time period exceeds 24 hours, the web browser might close. For per-minute data granularity, use the following calculation as a guideline: *number-of-days* x *number-of-hosts-selected-in-the-hosts-table* x *number-of-charts-to-display* ≤ 200

5. Identify the hosts that are increasing the load on the storage system based on IOPS or Data Transfer Rate metrics.

   To create a chart for IOPS or Data Transfer Rate metrics:

   a. In the **Hosts** table, sort the hosts by IOPS, or Data Transfer Rate metrics.
   b. Select the hosts with high IOPS or Data Transfer Rate.
   c. Select **New Analysis View**, and then click **Show Chart**.

      The data points are plotted for the default time period. The default time period is 3 to 4 hours.
   d. In the performance chart, you can view how the number of I/O operations or Data Transfer Rate varies. You can also determine the frequency of the peak and non-peak workload periods.
   e. (Optional) Change the time period and repeat the analysis to see how the workload changes over time and how it affects the overall performance of the host.

      To change the time range, select the start and end dates, select **Data Granularity**, specify the number of minutes, hours, or days, and then click **Apply**.

      You can also customize the time period settings by clicking the **Visual Settings** button, which lets you adjust the time range using the slider. In order for you to change the time range using the slider, one or more performance charts must be displayed in an analysis view.
6. Compare the performance charts to determine the hosts for which IOPS or Data Transfer Rate changes when the response time changes. Analyzing the performance metrics helps you identify the hosts that are increasing the workload on the storage system. If needed, you can export the analysis results to a PDF file.
7. (Optional)To analyze the performance of the volumes that are connected to a host, use the **Identify Performance Problems** wizard.

   a. On the **Analytics** tab, expand the **Hosts** tree, and select a host OS.
   b. In the hosts list, select the target host and click **Identify Performance Problems**, or click the host name link to select the target volumes and then click **Identify Performance Problems**.

      Follow the wizard prompts to continue the analysis.
8. (Optional) To improve the host application performance and balance I/O workload, allocate additional storage resources or distribute workload equally among existing storage resources.

Related concepts

- About analyzing storage system performance on page 526
- Allocating storage on page 141

Related tasks

- Identifying storage system performance problems on page 530
- Analyzing and load balancing MP Blades or MP Units on page 532

Metrics of volumes of VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, VSP F1500, Virtual Storage Platform, and HUS VM (Identify Performance Problems wizard)

Volume metrics for enterprise storage systems, as shown in the performance problem wizard, can help identify performance issues.

For each item used to analyze performance displayed in the Identify Performance Problems wizard, there are guideline values for evaluation if VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, VSP F1500, Virtual Storage Platform or Unified Storage VM volumes (basic volumes, DP volumes, external volumes) are used.

When analyzing DP volumes, all caches, MP Blades (VSP G1000, VSP G1500, VSP F1500, Virtual Storage Platform) and MP Units (VSP Fx00 models, VSP Gx00 models, Unified Storage VM) that are related to DP volumes are targeted. The target parity groups are those to which DP pool volumes that compose a DP pool belong. When analyzing external volumes, both the internal storage system and the external storage system are targeted.

In the Identify Performance Problems wizard, the status is judged as Caution when items exceed the following thresholds.

Table 23 Volume metrics for VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, VSP F1500, Virtual Storage Platform and Unified Storage VM (Identify Performance Problems wizard)

Item	Metric	Threshold	Notes
Port	Busy Rate	Caution: ≧ 100%	• With the default threshold value settings, the port status will not become [Caution].
Cache (CLPR by MP Blade/ Unit)	Write Pending Rate	Caution: > 30%	• Up to 30% is considered normal. • Frequent or sustained episodes of 40% deserve attention. • Frequent or sustained episodes to 50% require prompt remediation. • At 70%, inflow control is enabled which is disruptive. • High write pending is a sign of back end congestion.
MP Blade/Unit	Busy Rate	Caution: ≧ 40%	• -
Parity Group	Busy Rate	Caution: ≧ 50%	• Parity Group utilization should be less than 50%.

Item	Metric	Threshold	Notes
	Read Hit	Caution: ≦ 25%	• -

Tip: If you want to view detailed performance information, start Performance Reporter. By doing so, you can perform an analysis per minute or a long-term trend analysis for more performance evaluation items.

Related concepts

- About analyzing storage system performance on page 526

Metrics of volumes of Universal Storage Platform V/VM (Identify Performance Problems wizard)

For each item used to analyze performance displayed in the Identify Performance Problems wizard, there are guideline values for evaluation if Universal Storage Platform V/VM, basic volumes, DP volumes, or external volumes are used.

When analyzing DP volumes, all caches and CHPs (Channel Processors) that are related to DP pool volumes are targeted. The target parity groups and DKPs (Disk Processors) are those that are related to DP pool volumes that compose the DP pool. When analyzing external volumes, both the internal storage system and the external storage system are targeted.

In the Identify Performance Problems wizard, the status is judged as Caution when items exceed the following thresholds.

Table 24 Volume metrics for USP V/VM (Identify Performance Problems wizard)

Item	Metric	Threshold	Notes
Cache(CLPR)	Write Pending Rate	Caution: > 30%	• Up to 30% is considered normal. • Frequent or sustained episodes of 40% deserve attention. • Frequent or sustained episodes to 50% require prompt remediation. • At 70%, inflow control is enabled which is disruptive. • High write pending is a sign of back end congestion.
	Side File Usage Rate	Caution: ≧ 10%	• This is only applicable for TrueCopy Asynchronous. • Little, if any, should be reported under normal conditions. High side file usage is a sign of downstream congestion.
	Write Pending Rate + Side File Usage Rate	Caution: ≧ 70%	• It is very important to look at both CWP and SF together since they share the same cache pool and cannot exceed 70%. • The closer CWP or SF, or CWP+SF is to 70%, the closer the storage system is to major performance problems.

Item	Metric	Threshold	Notes
			• When the write pending reaches 70%, the storage system starts to destage write data until the write pending usage reaches 30%, which will degrade performance.
Channel Processor	Busy Rate	Caution: ≧ 45%	• 80% approximate maximum.
Parity Group	Busy Rate	Caution: ≧ 50%	• Parity Group utilization should be less than 50%.
	Read Hit	Caution: ≦ 25%	• -
Disk Processor	Busy Rate	Caution: ≧ 40%	• -

Tip: If you want to view detailed performance information, start Performance Reporter. By doing so, you can perform an analysis per minute or a long-term trend analysis for more performance evaluation items.

Related concepts

- About analyzing storage system performance on page 526

Metrics of volumes of midrange storage systems (Identify Performance Problems wizard)

For each item used to analyze performance displayed in the Identify Performance Problems wizard, there are guideline values for evaluation if volumes (basic volumes, DP volumes) of HUS 100, AMS 2000, SMS, or AMS/WMS are used.

When analyzing DP volumes, all processors and caches that are related to DP volumes are targeted. As for parity groups, those to which DP pool volumes, composing a DP pool, belong are targeted. The target drives (physical disks) are all those that compose a DP pool.

In the Identify Performance Problems wizard, the status is judged as Caution when items exceed the following thresholds.

Table 25 Volume metrics for midrange storage (Identify Performance Problems wizard)

Item	Metric	Threshold	Notes
Processor	Busy Rate	Caution: ≧ 80%	• Values over 66% are likely to be accompanied by response time increases due to contention. • Values over 50% will mean that the storage system will not be able to accommodate a controller failure without a reduction in workload capacity. • There are cases of resolving the problem by distributing the workload to other CPUs, or reducing

Item	Metric	Threshold	Notes
			the workload of such processes as TrueCopy, Copy-on-Write Snapshot, and ShadowImage.
Cache (CLPR)	Write Pending Rate on Controller 1, Write Pending Rate on Controller 0	Caution: > 25%	• The maximum Write Pending level per Partition is 70%. • There are cases of resolving the problem by distributing the workload to other controllers, adding to the cache, using high-speed drives, adding drives to parity groups, adding parity groups, or switching to RAID 1+0.
Parity Group	Write Hit	Caution: < 100%	• -
Drive	Operating Rate on Controller 1, Operating Rate on Controller 0	Caution: ≧ 70%	• Values over 66% will result in response deterioration due to contention.

Tip: If you want to view detailed performance information, start Performance Reporter. By doing so, you can perform an analysis per minute or a long-term trend analysis for more performance evaluation items.

Related concepts

- About analyzing storage system performance on page 526

Creating periodic health check reports for storage systems performance analysis

You can create storage system health check (performance) reports, and control scheduling, the data collection time period, and data interval settings. You can review the health check reports when completed, and later by reviewing health check history.

Note: When you complete health checks periodically or immediately, do not perform the following operations simultaneously. If you do, the health checks (or subsequent operations) take more time to finish:

- Updating storage system configurations or information about storage systems.
- Periodically updating configuration information or performance information that is required for Universal Replicator performance analysis.

Additionally, if you are using the UR Performance Analysis window of the Replication tab for Universal Replicator performance analysis, and a periodic refresh schedule has been configured for data collection, make sure that

storage system health checks (performance reports) run at a different time to support the distribution of processing loads.

Before you begin

- Install Tuning Manager and link it to HCS.
- Register the Tuning Manager license.
- Import the report definition file in Tuning Manager before analyzing performance with the Analytics tab (ensure that you delete the report definition file for the previous version before importing the latest report definition file).
- To receive the results of Health Check execution by email, you must specify the SMTP server and edit the `server.properties` file of the Device Manager server.
- For more details on these tasks, see the *Hitachi Command Suite Administrator Guide*.

Procedure

1. On the **Analytics** tab, select **Storage Systems**.
2. Select one or more storage systems, and click **Create Schedule**.
3. Specify the required items, and run the task.
4. After verifying that the task is complete, check for any performance problems. If there is a problem, check the Health Check report of the corresponding storage system. In accordance with the specified threshold values, the background of the displayed analysis result chart is colored for each threshold range. The Caution threshold range is displayed in yellow. The Warning threshold range is displayed in pink.
5. Verify the problem areas from the report information, and then analyze and identify the problems.



Tip: Observe the following guidelines when analyzing the report information:

- If you want to analyze the performance of storage systems immediately, select the target storage systems, and then click **Run Health Check**. You can view the execution results in the **Tasks & Alerts** tab.
- If you want to delete a schedule, from **Schedule**, select the target schedule from the **Health Check Schedules** list, and then click **Delete Schedule**.
- If you want to view detailed performance information, start Tuning Manager from the **Tools** menu. By displaying Performance Reporter reports from Tuning Manager, you can analyze performance based on a per-minute basis or on a long-term trend basis. For information about how to use Performance Reporter, see the *Hitachi Command Suite Tuning Manager User Guide*.

- If you want to view performance information about the resources associated with a pool, select the target storage system, then on the **DP Vols** tab, or the **Parity Groups** tab, display the Historical Report by clicking the **View Performance Reports** button. For information about changing the report window settings for the Historical Report, see the *Hitachi Command Suite Tuning Manager User Guide*.

Related concepts

- About analyzing storage system performance on page 526

Related tasks

- Editing periodic health check schedule on page 542
- Exporting health check reports on page 542
- Viewing and deleting health check reports on page 543

Editing periodic health check schedule

You can change a previously created storage system health check schedule.

Procedure

1. From the **Analytics** tab, select **Schedule**.
2. Select the schedule to change, and click **Edit Schedule**.
3. Specify the required items, and then execute the task.

Result

Display the window showing the target schedule in detail, and confirm that the schedule has been changed.

Tip: If you are using the UR Performance Analysis window of the Replication tab, and a periodic refresh schedule has been configured for data collection, make sure that storage system health checks (performance reports) run at a different time to ensure the distribution of processing loads.

Related concepts

- About analyzing storage system performance on page 526

Related tasks

- Creating periodic health check reports for storage systems performance analysis on page 540

Exporting health check reports

You can export health check reports to a PDF or CSV file for reporting purposes.

Procedure

1. From the **Analytics** tab, select **Report History**.
2. From the **Reports** list, select the report to export, and click **Export Reports**.
3. Select the required items and click **Submit**.
4. Select the export location for the health check report.

Result

A CSV file or PDF file is saved in the specified location.

Related concepts

- About analyzing storage system performance on page 526

Related tasks

- Viewing and deleting health check reports on page 543

Viewing and deleting health check reports

Previously created health check reports can be reviewed to check for and resolve performance issues. The report list can display up to 500 health check reports. If the limit is exceeded, health check reports are automatically deleted starting with the oldest, first.

Procedure

1. From the **Analytics** tab, select **Report History**.
2. From the **Reports** list, select the report to review, and click **View Report** of the **Health Status** column.
3. Review the displayed health check report.

   In accordance with the specified threshold values, the background of the displayed analysis result chart is colored for each threshold range. The Caution threshold range is displayed in yellow. The Warning threshold range is displayed in pink.

Result

Based on the health check report, resolve performance issues if necessary.

Tip: To delete reports that are no longer needed, select the reports to be deleted, and click Delete Reports button.

Related concepts

- About analyzing storage system performance on page 526

Related tasks

- Creating periodic health check reports for storage systems performance analysis on page 540

- Exporting health check reports on page 542

Metrics of resources of VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, VSP F1500, Virtual Storage Platform, and HUS VM (Health Check Report)

Volume metrics for enterprise storage systems, as shown in the health check report, can help identify performance issues.

The following provides the values for evaluating items that are used by Health Check to verify whether there is a performance problem for VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, VSP F1500, Virtual Storage Platform or HUS VM.

During Health Check, if items exceed the following thresholds, the status is judged as Caution or Warning.

Table 26 Resource metrics for VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, VSP F1500, Virtual Storage Platform and HUS VM (Health Check Report)

Item	Metric	Threshold	Notes
Cache Total (entire storage system), Cache (CLPR), Cache (CLPR by MP Blade/Unit)	Write Pending Rate	Caution: > 30% Warning: > 60%	• Up to 30% is considered normal. • Frequent or sustained episodes of 40% deserve attention. • Frequent or sustained episodes to 50% require prompt remediation. • At 70%, inflow control is enabled which is disruptive. • High write pending is a sign of back end congestion.
MP Blade/Unit	Busy Rate	Caution: ≧ 40% Warning: ≧ 80%	• -
Parity Group	Max Busy Rate	Caution: ≧ 50%, Warning: ≧80%	• -

Tip: If you want to view detailed performance information, start Performance Reporter. By doing so, you can perform an analysis per minute or a long-term trend analysis for more performance evaluation items.

Related concepts

- About analyzing storage system performance on page 526

Metrics of resources of Universal Storage Platform V/VM (Health Check Report)

The following provides the values for evaluating items that are used by Health Check to verify whether there is a performance problem for the Universal Storage Platform V/VM.

During Health Check, if items exceed the following thresholds, the status is judged as Caution or Warning.

Table 27 Resource metrics of USP V/VM (Health Check Report)

Item	Metric	Threshold	Notes
Cache Total (entire storage system), Cache (CLPR)	Write Pending Rate	Caution: > 30% Warning: > 60%	• Up to 30% is considered normal. • Frequent or sustained episodes of 40% deserve attention. • Frequent or sustained episodes to 50% require prompt remediation. • At 70%, inflow control is enabled which is disruptive. • High write pending is a sign of back end congestion.
	Side File Usage Rate	Caution: ≧ 10% Warning: ≧ 20%	• This is only applicable for TrueCopy Asynchronous. • Little, if any, should be reported under normal conditions. High side file usage is a sign of downstream congestion.
Channel Processor	Busy Rate	Caution: ≧ 45% Warning: ≧ 80%	• 80% approximate maximum.
Parity Group	Max Busy Rate	Caution: ≧ 50%, Warning: ≧80%	• -
Disk Processor	Busy Rate	Caution: ≧ 40% Warning: ≧ 80%	• -

Tip: If you want to view detailed performance information, start Performance Reporter. By doing so, you can perform an analysis per minute or a long-term trend analysis for more performance evaluation items.

Related concepts

- About analyzing storage system performance on page 526

Metrics of resources of midrange storage systems (Health Check Report)

The following provides the values for evaluating items that are used by Health Check to verify whether there is a performance problem for the HUS 100, AMS 2000, SMS, and AMS/WMS platforms.

During Health Check, if items exceed the following thresholds, the status is judged as Caution or Warning.

Table 28 Resource metrics of midrange storage (Health Check Report)

Item	Metric	Threshold	Notes
Cache (CLPR)	Write Pending Rate on Controller 1, Write Pending Rate on Controller 0	Caution: > 25% Warning: > 50%	• The maximum Write Pending level per Partition is 70%. • There are cases of resolving the problem by distributing the workload to other controllers, adding to the cache, using high-speed drives, adding drives to parity groups, adding parity groups, or switching to RAID 1+0.
Processor	Busy Rate	Caution: ≧ 80% Warning: ≧ 90%	• -

Tip: If you want to view detailed performance information, start Performance Reporter. By doing so, you can perform an analysis per minute or a long-term trend analysis for more performance evaluation items.

Related concepts

- About analyzing storage system performance on page 526

Changing performance threshold settings to correctly analyze storage system performance

The recommended threshold values for analyzing storage system performance are set in advance. If the recommended values are not suitable, you can change the threshold values. Use the metrics of volumes of each storage system as a reference to change the threshold.

There are four types of thresholds:

- Thresholds that are used only in the Identify Performance Problems wizard.
- Thresholds that are used only in the Health Check window.
- Thresholds that are used only in the Analyze MP Blades/Units window.
- Thresholds that are common to the Identify Performance Problems wizard, the Analyze MP Blades/Units window and the Health Check window.

Note: When you change a threshold common to Identify Performance Problems, Analyze MP Blades/Units, and Health Check, the change applies to all analysis tools in their respective windows.

Tip: For the specified threshold values, the background of the displayed analysis result chart is colored for each threshold range. The Caution threshold range displays in yellow. The Warning threshold range displays in pink.

Procedure

1. On the **Analytics** tab, in **General Tasks**, select **Performance Threshold Settings**.
2. Select the series of the target storage system, and then use the metrics of volumes as a reference to change the threshold value for each metric you want to change.
3. Change the threshold settings.
   a. Select the target storage system.
   b. Select the analysis tool whose metrics you want to change (for example Analyze MP Blades/Units) or select **All** to display and change all available metrics.

      Tip: Metrics might be deactivated based on your analysis tool selection.

   c. Change the caution and warning (Health Check only) thresholds to meet your requirements.
   d. Click **OK**.

Result

Metrics are changed for one or more analysis tools.

Related concepts

- About analyzing storage system performance on page 526

Related tasks

- Identifying storage system performance problems on page 530
- Analyzing and load balancing MP Blades or MP Units on page 532
- Creating periodic health check reports for storage systems performance analysis on page 540

Related references

- Metrics of volumes of VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, VSP F1500, Virtual Storage Platform, and HUS VM (Identify Performance Problems wizard) on page 537
- Metrics of volumes of Universal Storage Platform V/VM (Identify Performance Problems wizard) on page 538
- Metrics of volumes of midrange storage systems (Identify Performance Problems wizard) on page 539
- Metrics of resources of VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, VSP F1500, Virtual Storage Platform, and HUS VM (Health Check Report) on page 544
- Metrics of resources of Universal Storage Platform V/VM (Health Check Report) on page 545
- Metrics of resources of midrange storage systems (Health Check Report) on page 546

Configuring the Performance Reporter display

After activating Performance Reporter, you can analyze detailed storage system information based on Identify Performance Problems results or Health Check report results.

The following table lists the configurable resources for displaying performance and status information based on storage system, selected resources, and the available Identify Performance Problems or Health Check information.

Table 29 Performance Reporter configurable resources for display

Storage System	Resource	Value
VSP Fx00 models, VSP Gx00 models VSP G1000, VSP G1500, VSP F1500 Virtual Storage Platform Unified Storage VM	LDEV Number	Value of the Volume column
	CLPR Number	Value of the CLPR column
	Adaptor ID	Value of the MP Blade/Unit column
	MP Blade ID	Value of the MP Blade/Unit column
	Processor ID (for MP Blade or MP Unit)	"*_Total*": Default performance value displayed when you start Performance Reporter.
	Processor ID (for MP)	The numeric value after the last hyphen (-) in the MP column. For example, If the value of the MP column is MPU-L0-00 , specify 00.
	RAID Group Number	Character string with asterisk (*) wildcard characters before and after the value of the Parity Group column. For example, if the value of the Parity Group column is 1-3, specify *1-3*. If you want to further narrow down the displayed information, change the wildcard character specification.
	Port Name	Character string of the Port column

Storage System	Resource	Value
	Pool ID	Numerical value of the Pool column without "DP" at the beginning. For example, if the value of the Pool column is "DP 10", specify 10.
Universal Storage Platform V/VM	LDEV Number	Value of the Volume column
	CLPR Number	Value of the CLPR column
	Adaptor ID	Character string before the second hyphen in the value of the Processor column. For example, if the value for the Processor column is CHA-1FL-7, specify CHA-1FL.
	Processor ID	Character string after the second hyphen in the value of the Processor column. For example, if the value for the Processor column is CHA-1FL-7, specify 7.
	RAID Group Number	Character string before the second hyphen in the value of the Parity Group column, and with asterisk (*) wildcard characters before and after. For example, if the value of the Parity Group column is 1-2-1, specify *1-2*. If you want to further narrow down the displayed information, change the wildcard character specification.
	Port Name	Character string of the Port column.
	Pool ID	Numerical value of the Pool column without the DP at the beginning. For example, if the Pool column value is DP 10, specify 10.
HUS 100 AMS 2000 SMS AMS/WMS	LDEV Number	Value of the Volume column. If the value of the Volume column starts with 0, such as 01, specify the numerical value without the 0 at the beginning.
	CLPR Number	Value of the CLPR column
	Controller	Controller value displayed in the chart title
	Processor ID	Character string of the Processor column
	RAID Group Number	Value of the Parity Group column
	Unit Number	Numerical value before the hyphen in the value of the Drive column. For example, if the value of the Drive column is 0-7, specify 0.
	HDU Number	Numerical value after the hyphen in the value of the Drive column. For example, if the value of the Drive column is 0-7, specify 7.
	Port Name	Character string without the second character, T, in the Port column. For example, if the value of the Port column is CTL0-B, specify CL0-B.

Related concepts

- About analyzing storage system performance on page 526

Related tasks

- Identifying storage system performance problems on page 530
- Analyzing and load balancing MP Blades or MP Units on page 532
- Creating periodic health check reports for storage systems performance analysis on page 540

Related references

- Metrics of volumes of VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, VSP F1500, Virtual Storage Platform, and HUS VM (Identify Performance Problems wizard) on page 537
- Metrics of volumes of Universal Storage Platform V/VM (Identify Performance Problems wizard) on page 538
- Metrics of volumes of midrange storage systems (Identify Performance Problems wizard) on page 539
- Metrics of resources of VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, VSP F1500, Virtual Storage Platform, and HUS VM (Health Check Report) on page 544
- Metrics of resources of Universal Storage Platform V/VM (Health Check Report) on page 545
- Metrics of resources of midrange storage systems (Health Check Report) on page 546
- Performance Reporter and Analytics tab display differences on page 550

Performance Reporter and Analytics tab display differences

The displayed report information differs between the Analytics tab and Performance Reporter.

The following table describes how the displayed information of the Analytics tab corresponds to that of Performance Reporter.

Table 30 Analytics tab and Performance Reporter display differences

Function	Resource	Metric	Record Name(Record ID)	Field Name
Identify Performance Problems	Volume[1]	IOPS	VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, VSP F1500, Virtual Storage Platform, Unified Storage VM, Universal Storage Platform V/VM: Logical Device Summary (PI_LDS, PI_LDS1, PI_LDS2, PI_LDS3) HUS 100, AMS 2000, SMS, AMS/ WMS: Logical Device Summary (PI_LDS)	Read I/O /sec + Write I/O /sec
	Pool[1]	Response Time	Pool Summary (PI_PLS)	((Read-Response-Rate x Read I/O / sec) + (Write-Response-Rate x Write I/O /sec)) / (Read I/O /sec + Write I/O /sec)

Function	Resource	Metric	Record Name(Record ID)	Field Name
		IOPS	Pool Summary (PI_PLS)	Read I/O /sec + Write I/O /sec
	Port	Busy Rate	Port Summary (PI_PTS)	Avg Xfer/sec / actual-port-speed x 100
	Cache(CLPR)	Write Pending Rate + Side File Usage Rate	CLPR Summary (PI_CLPS)	Cache Write Pending Usage % + Cache Side File Usage %
	Parity Group	Read Hit	RAID Group Summary (PI_RGS)	Read Hit I/O Count / Read I/O Count x 100
	Parity Group	Write Hit	RAID Group Summary (PI_RGS)	Write Hit I/O Count / Write I/O Count x 100
Health Check	System Total FE IOPS	IOPS	Logical Device Aggregation (PI_LDA)	Read I/O /sec + Write I/O /sec
	System Total FE MB/s	MB/s	Logical Device Aggregation (PI_LDA)	Read Xfer /sec + Write Xfer /sec
	System Total BE IOPS	IOPS	Logical Device Aggregation (PI_LDA)	VSP Fx00 models, VSP Gx00 models, VSP G1000, VSP G1500, VSP F1500, Virtual Storage Platform, HUS VM, Universal Storage Platform V/VM: Read I/O /sec x (1 - Read Hit %) + Write I/O /sec HUS 100, AMS 2000, SMS, AMS/WMS: Read I/O /sec x (1 - Read Hit %) + Write I/O /sec x (1 - Write Hit %)
	Busy PG by FE IOPS	IOPS	RAID Group Summary (PI_RGS)	Read I/O /sec + Write I/O /sec
	Busy PG by FE MB/s	MB/s	RAID Group Summary (PI_RGS)	Read Xfer /sec + Write Xfer /sec
	Busy Pools by IOPS	IOPS	Pool Summary (PI_PLS)	Read I/O /sec + Write I/O /sec
	Busy Pools by MB/s	MB/s	Pool Summary (PI_PLS)	Read Xfer /sec + Write Xfer /sec
	Busy PG by BE IOPS	IOPS	RAID Group Summary (PI_RGS)	SMS, AMS/WMS: Read I/O /sec x (1 - Read Hit %) + Write I/O /sec x (1 - Write Hit %)

Function	Resource	Metric	Record Name(Record ID)	Field Name
[1] Response Time for volumes and pools are displayed in *microseconds* in the Performance Reporter, and is displayed in *milliseconds* in the Analytics tab.				

Related references

- Configuring the Performance Reporter display on page 548

Analyzing Universal Replicator performance

This module describes procedures for analyzing Universal Replicator performance to identify potential problems.

About analyzing Universal Replicator performance

The UR Performance Analysis window of the Replication tab provides information for analyzing performance problems with data transfers between the primary and secondary storage systems.

Universal Replicator (UR) asynchronously transfers data to the remote site. Delay times occur and differ depending on the data transfer timing. This delay is known as *C/T delta* and is an indicator for Recovery Point Objective (RPO).

You can use the UR Performance Analysis window to:

- Identify instances where the C/T delta threshold was exceeded.
  View the top five copy groups and their C/T delta rates in a chart. You can quickly identify copy groups that consistently exceed the maximum write time delay (C/T delta threshold). (You can also check the status of copy groups in the Alerts table and edit the alert settings or the C/T delta threshold if desired.)
- Analyze the C/T delta threshold against the performance of primary, secondary, and network resources. The analysis process supports two modes:
  - **Wizard mode:** a step-by-step guide that compares trends and helps users identify the cause of the problem.
  - **Advanced mode:** a selection of charts that lets advanced users correlate multiple trends and identify the problematic resource.

Related tasks

- Monitoring Universal Replicator performance on page 553
- Configuring network bandwidth for analysis on page 555
- Analyzing Universal Replicator performance in wizard mode on page 559
- Analyzing Universal Replicator performance in advanced mode on page 560

Related references

- Prerequisites for analyzing Universal Replicator performance on page 553
- Metrics for Universal Replicator performance analysis on page 556

Prerequisites for analyzing Universal Replicator performance

Before attempting analysis, verify the site configuration, necessary software, and user access controls.

You must satisfy the following requirements before using the UR Performance Analysis window:

- Configure the alert setting for C/T delta in the Replication Management window.
- Ensure that your site configuration meets all the configuration requirements for Universal Replicator performance analysis described in the *Hitachi Command Suite Administrator Guide*.
- Ensure that Hitachi Tuning Manager and Replication Manager are installed and linked to HCS.
- Ensure that users have the necessary permissions (roles and resource access) for the UR Performance Analysis window of the Replication tab described in Required roles and resource groups by function on page 114.

Note: The performance analysis function supports copy groups consisting of open volumes. Copy groups that are defined by device group, use multiple journal groups in a Universal Replicator MxN configuration, or use virtual command devices, are not supported.

Related concepts

- About analyzing Universal Replicator performance on page 552

Related tasks

- Setting alerts for a replication configuration on page 314
- Monitoring Universal Replicator performance on page 553

Related references

- Required roles and resource groups by function on page 114

Monitoring Universal Replicator performance

You can monitor Universal Replicator performance (C/T delta) from the UR Performance Analysis window. If you see problems, you can then perform an analysis to discover the root cause.

Tip: By default, performance data is collected every four hours, and configuration information (including storage systems, copy pairs, and pair management servers) is refreshed every 24 hours. The Refresh Performance

Data button allows you to do a manual refresh of the performance data, and displays an option to refresh the configuration as well (Also refresh copy group information).

Procedure

1. Open the **Replication** tab, **Performance Analysis**.
2. Examine the C/T delta trend in the UR Performance Analysis window:
   - The five copy groups with the largest percentage of C/T delta threshold are displayed.
   - The text in the bar graph displays the actual C/T delta and threshold values.
   - The default Time Period is one day. Use the range buttons to adjust the view.
   - The **Copy Groups** tab includes a table with copy group information and key usage statistics.
   - For reference, the **Alerts** tab includes a summary of alerts that are in the Replication Management window.
3. If the C/T delta threshold for a copy group is exceeded, you can perform an analysis:
   - In *advanced mode*, you can perform the analysis yourself by viewing a series of graphs to determine the root cause of the problem.
   - In *wizard mode*, you are guided through a selection of key metrics and presented with a summary of possible causes and solutions.

Caution: The following operations will cause previous C/T delta values and JNLG usage rates not to be displayed:
- Deleting copy pair configuration definition files
- Making changes to the pair management server that manages a target copy group, including changing the WWN

Note: You cannot access previous C/T delta values during a takeover operation (disaster recovery) because the relationship between the primary and secondary volumes is temporarily reversed. When the copy group relationship is restored (takeback), the data can be accessed again. For more information about disaster recovery operations, see the *Hitachi Command Suite Replication Manager User Guide*.

Related concepts

- About analyzing Universal Replicator performance on page 552

Related tasks

- Analyzing Universal Replicator performance in advanced mode on page 560
- Analyzing Universal Replicator performance in wizard mode on page 559
- Configuring network bandwidth for analysis on page 555

Related references

- Prerequisites for analyzing Universal Replicator performance on page 553
- Metrics for Universal Replicator performance analysis on page 556

Configuring network bandwidth for analysis

The UR Performance Analysis function gives you the option of supplying values for the effective network bandwidth.

The *effective* network bandwidth is the actual speed at which data can be transmitted on a remote path based on your replication environment. Check your network and supply a proper bandwidth value for each path group.

Tip:

- In addition to appearing under General Tasks, more, this option is also available during analysis (wizard and advanced mode).
- Setting the bandwidth is not required. If values are not set, a warning is displayed that can be ignored.

Procedure

1. On the **Replication** tab, in **General Tasks**, click **more** and select **Configure Thresholds for Network Analysis**.
2. Supply one value for all remote path groups or select each individually.

Related concepts

- About analyzing Universal Replicator performance on page 552

Related tasks

- Monitoring Universal Replicator performance on page 553

Related references

- Prerequisites for analyzing Universal Replicator performance on page 553
- Metrics for Universal Replicator performance analysis on page 556

Configuring metric thresholds for M-JNL analysis

The UR Performance Analysis window includes an option to set thresholds for the following metrics:

- Host Write Transfer Rate to M-JNL
- M-JNL Async Transfer Rate to RCU

- Host Write IOPS to M-JNL

The threshold values are used to plot a horizontal lines in graphs indicating where the limit has been exceeded. Although defaults are defined, the values should be based on your replication environment.

Tip: In addition to appearing under General Tasks, more, this option is also available during analysis (wizard and advanced mode).

Procedure

1. On the **Replication** tab, in **General Tasks**, click **more** and select **Configure Thresholds for M-JNL Analysis**.
2. Select a metric using the M-JNL Threshold buttons.
3. Choose **All** copy groups or **Select**. You can select from the Available Copy Groups and use the **Add/Remove** buttons to modify the list of Selected Copy Groups.
4. Enter a new value (or revert to the default) and click **Apply**.
5. Repeat process for each of the desired metrics. Click **Close** when you are finished.

Metrics for Universal Replicator performance analysis

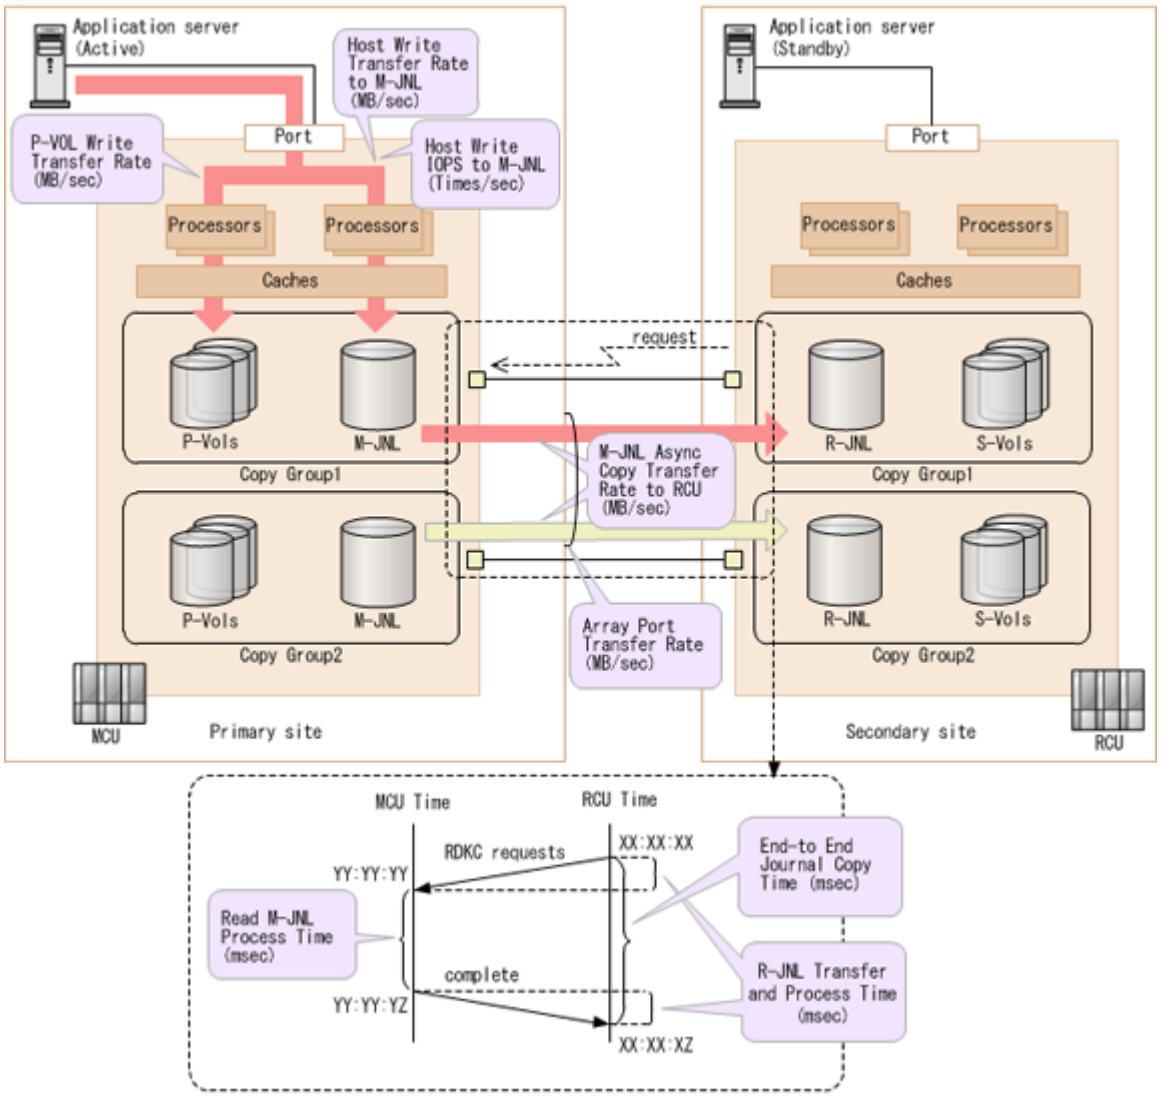
The following table lists the metrics used to analyze Universal Replicator performance. See the figure for an illustration of how each metric relates to the data flow.

Table 31 Universal Replicator performance metrics

Metric	Description	Threshold value / notes
Cache Write Pending Rate	Write pending rate for the cache used by the MPB.	Caution: > 30%
Processor Busy Rate	Operating rate of the MPB.	Caution: > 40%
Journal (JNL) Usage Rate	Usage rate for the journal.	This metric does not directly lead to C/T delta degradation. The system status can be determined by checking JNLG usage. Secondary journal usage available in advanced mode only.
Array Port Transfer Rate	Amount of data transmitted from the MCU to the RCU per remote path.	Caution: > network bandwidth value See Configuring network bandwidth for analysis on page 555.
Host Write IOPS to M-JNL	Average amount of data input from the host to the primary JNL per copy group (MBs per second)	Threshold value for graphs can be set as described in Configuring metric thresholds for M-JNL analysis on page 555. Default: 12,000.
Host Write Transfer Rate to M-JNL	Frequency of write operations for the primary JNL per unit of time	Threshold value for graphs can be set as described in Configuring metric thresholds for M-JNL analysis on page 555. Default: 250.

Metric	Description	Threshold value / notes
Host Write Block Size to M-JNL	Amount of data input from the host to the primary JNL for each write operation.	Available in advanced mode only.
P-VOL Write Transfer Rate	Amount of data transmitted from the host or host group per P-VOL (MBs per second).	Available in advanced mode only. See About the P-VOL Write Transfer Rate metric on page 562 for additional information.
M-JNL Async Transfer Rate to RCU	Average rate of transmission from M-JNL to R-JNL (MBs per second).	Threshold value for graphs can be set as described in Configuring metric thresholds for M-JNL analysis on page 555. Default: 250.
Read M-JNL Process Time	Interval between when the MCU receives the RDJNL command and the MCU sends a data transmission completion command to the RCU (msec).	Available in advanced mode only.
End-to-End Journal Copy Time	Interval between when the RCU sends the RDJNL command to the MCU and the RCU receives a data transmission completion command (msec).	Available in advanced mode only.
R-JNL Transfer and Process Time	The time lag in network transmission between M-JNL and R-JNL.	This value represents the difference between the End-to-End Journal Copy Time and Read M-JNL Process Time. (If one of these values is not available, the R-JNL Transfer and Process Time value is undefined.)

The following figure depicts points in the data flow associated with key metrics. This is not intended as a master reference for all metrics, but serves as an aid for performing an analysis in advanced mode. For example, the expanded portion shows the relationship between the Read M-JNL Process Time, R-JNL Transfer and Process Time, and the End-to-End Journal Copy Time.



Related concepts

- About analyzing Universal Replicator performance on page 552

Related tasks

- Monitoring Universal Replicator performance on page 553
- Configuring network bandwidth for analysis on page 555

Related references

- Prerequisites for analyzing Universal Replicator performance on page 553

Analyzing Universal Replicator performance in wizard mode

In wizard mode, the GUI guides you through the process of examining resource metrics and suggests possible solutions.

Procedure

1. Open the **Replication** tab, **Performance Analysis**.
2. On the **Copy Groups** tab, select a copy group and click **Analyze UR Performance (Wizard Mode)**.
3. Read the **Introduction** and click **Next**.
4. Using the slider below the graph, choose a **Time Period** as instructed and then click **Next**.

   Note: If you see cache or processor resources listed that are not associated with the copy group you are analyzing, execute **Refresh Storage System** on the **Administration** tab and then, select **Refresh Performance Data** on the **Replication** tab with the **Also refresh copy group configuration** check box selected.

   - By default, Replication Manager polls information sources for copy pair status every 5 minutes. Depending on the data interval you set, this can affect how the display is refreshed in the UR Performance Analysis window.
   - By default, new performance data is collected through Tuning Manager every four hours. A gray zone in the graph indicates uncollected data. If you click **Refresh Performance Data** at this point, the wizard closes and you restart this procedure when the task is complete.
5. On the displayed pages, examine each metric and use the check box to mark any resources that exceed the threshold. (You can select **Mark resources that exceed pre-set threshold** to select them all.) Click **Next** to proceed.
6. The **Results Report** page includes the detected issues and the possible causes and solutions based on your selections. Click the **Show Detail** link to display the resources and the graphed data.

Result

The solutions include instructions to check each resource to confirm the problem and apply the suggested solution.

Related concepts

- About analyzing Universal Replicator performance on page 552
- About refreshing UR Performance Analysis data on page 563

Related tasks

- Analyzing Universal Replicator performance in advanced mode on page 560
- Exporting Universal Replicator performance reports on page 561

Related references

- Prerequisites for analyzing Universal Replicator performance on page 553
- Metrics for Universal Replicator performance analysis on page 556

Analyzing Universal Replicator performance in advanced mode

In advanced mode, UR Performance Analysis creates graphs that are based on a selection of metrics to highlight possible performance problems.

Note: If you see cache or processor resources listed that are not associated with the copy group you are analyzing, execute Refresh Storage System on the Administration tab and then, select Refresh Performance Data on the Replication tab with the Also refresh copy group configuration check box selected.

Procedure

1. Open the **Replication** tab, **Performance Analysis**.
2. On the **Copy Groups** tab, select a copy group and click **Analyze UR Performance (Advanced Mode)**.

Result

The default dashboard includes four panes that plot metrics for **S-VOL**, **M-JNL**, **R-JNL**, and **Network**. Each panel has a Report list to choose metrics to plot in a graph. Use the maximize control to view a single pane.

Tip: The M-JNL pane includes reports featuring the *P-VOL Write Transfer Rate*, a metric that displays the amount of data transmitted from host or host group per P-VOL.

You can customize the display in several ways:

- Time Settings adjusts the time period for all panes.
- View Settings selects the panes to display in the dashboard.
- Each report includes a list for selecting the metric, and Configure Chart customizes the display:
  - You can select an existing entry from the Charts list to modify (or remove). Clicking Add creates a chart.
  - Chart Settings sets the title, chooses metrics for the left and right axes, and the chart position.
- When a pane is maximized, Vertical Comparison adds a second graph.
- When the Network pane is maximized, the graph is maximized.

- If you click the Volumes Sharing Path Group button while the **Network** pane is maximized, a list of primary volumes in the copy group that uses the same remote path as the Universal Replicator copy group that is being analyzed is displayed in the Related Volumes window.

- You can also annotate a graph with *memos*, text windows that allow you to take notes.
  - Click Add Memo to start a new note. You can also add a note by double-clicking within a graph.
  - As with any window, you can move, minimize, or delete it using the controls in the title bar.
- Configure Thresholds for M-JNL Analysis changes the threshold value (horizontal red line) in the chart.

Tip: By default, Replication Manager polls information sources for copy pair status every 5 minutes. Depending on the data interval you set, this can affect how the display is refreshed in the UR Performance Analysis window.

Related concepts

- About analyzing Universal Replicator performance on page 552
- About the P-VOL Write Transfer Rate metric on page 562
- About refreshing UR Performance Analysis data on page 563

Related tasks

- Analyzing Universal Replicator performance in wizard mode on page 559
- Exporting Universal Replicator performance reports on page 561

Related references

- Metrics for Universal Replicator performance analysis on page 556
- Prerequisites for analyzing Universal Replicator performance on page 553
- Metrics for Universal Replicator performance analysis on page 556

Exporting Universal Replicator performance reports

You can export Universal Replicator performance data to a PDF file for reporting purposes.

The UR Performance Analysis window includes an export function to save output to a PDF file on the following panes:

- Analyze UR Performance (main screen)
- Results Report (wizard mode)
- Analysis Target (advanced mode)

You can also obtain UR performance data (in CSV form) using the following command line tools:

- `GetCopyGroup`
- `RefreshReplicationPerformanceData`
- `GetReplicationPerformanceData`

For more information, see the *Hitachi Command Suite CLI Reference Guide*.

Reporting tip

It is useful to save a C/T delta baseline so that you can later detect any gradual worsening of C/T delta performance over time. An accurate baseline requires stable host I/O. Wait for the completion of any settings and status changes that affect C/T delta (such as initial copies of newly added copy groups, or a copy group that shares a remote path). For best results, when using the Export function on the main screen, select a Time Period of one day (1d).

Procedure

1. If you are exporting data from the Analyze Universal Replicator performance main screen, first select a copy group from the **Copy Groups** tab.
2. Click **Export**.
   If you are exporting data from the main screen, you are prompted to select a Time Period.
3. For graphs in the advanced mode pane, any memos you created are included by default. To exclude your memos from the report, de-select **Display the memos on the graphs**.
4. Click **Submit**.
5. Select the export location for the report.
   The PDF file is saved in the specified location.

Related tasks

- Analyzing Universal Replicator performance in wizard mode on page 559
- Analyzing Universal Replicator performance in advanced mode on page 560

Related references

- Prerequisites for analyzing Universal Replicator performance on page 553

About the P-VOL Write Transfer Rate metric

The P-VOL Write Transfer Rate metric is used to calculate the amount of data written (in MB/sec) for each of the P-VOLs that comprise a copy group. Normally, the total of P-VOL Write Transfer Rate values in a copy group approximates the value of the Host Write Transfer Rate to M-JNL. The P-VOL Write Transfer Rate metric is available in the M-JNL pane of the advanced mode window.

The P-VOL Write Transfer Rate can be used to narrow down possible causes after identifying the copy group where C/T delta degradation has been observed. There are two basic scenarios:

- *P-VOL traffic may be dominating network bandwidth.*

For best results, first observe the Array Port Transfer Rate (amount of data output from a port) or the M-JNL Async Transfer Rate to RCU (amount of data output for each M-JNL). Once you have identified the specific copy group, you can use the P-VOL Write Transfer Rate to determine which application/host uses the constituent P-VOLs, or which host group has written the largest amount of data.

- *P-VOL traffic may be overloading storage resources.*
  If you observe the Host Write Transfer Rate to M-JNL increase during or before a C/T Delta increase, storage resources (such as processor, cache, or parity group) may be overtaxed. In this case, analyze the P-VOL Write Transfer Rate to identify the responsible host/application. You can then choose to tune the application or add resources.



Tip: When you select a report that includes the P-VOL Write Transfer Rate, the graph includes a filter that permits you to narrow down the number of P-VOLs to be plotted.

About refreshing UR Performance Analysis data

By default, Replication Manager polls information sources for copy pair status every 5 minutes. (This setting is part of the Replication Manager GUI, not the UR Performance Analysis window.) If you have increased this polling setting (refresh value) to be longer than the data interval in the UR Performance Analysis window, the plot in the graph is updated less often (regardless of the data interval setting). For example, a refresh value of 30 minutes updates the graph every half hour, even if the data interval is set to 10 minutes. Setting the refresh value to less than 10 minutes helps ensure the plot is updated every 10 minutes (assuming data interval of 10 minutes). However, you might also need to consider the load balance if you reduce the refresh value because (in some cases) refreshing the copy pair status will actually take longer than the setting indicates.

Managing alerts

This module describes how to view, confirm, and manage alerts.

About alerts

Hitachi Command Suite regularly monitors storage system (including NAS modules) and file server (NAS Platform v12.0 or later only) status.

Alerts are reported on the Dashboard, in the Storage System Alerts report. You can select the severity levels of alerts to be collected in this report by clicking the tool icon application pane.

Alert details can be confirmed on the Tasks & Alerts tab.

- Resolved alerts can be deleted manually from the list.

- The upper limit of the alerts that are managed by HCS is 500,000. When the number of alerts exceeds the upper limit, older alerts are deleted.

For VSP Fx00 models, VSP Gx00 models, you can also receive alerts through SNMP traps or email. For information about alerts for VSP Fx00 models, VSP Gx00 models and about how to configure SNMP trap and email notification settings, see the *System Administrator Guide* for your storage system.

Alerts can also be delivered in email when SNMP traps are received from storage systems or file servers (NAS Platform v12.0 or later, NAS Platform F, and Data Ingestor only). You can also configure VSP G1000, VSP G1500, or VSP F1500 to receive failure information (SIM) through syslog, SNMP trap, and email. For information about how to configure HCS and your storage to receive failure information, see the *System Administrator Guide* for your storage system.

Related tasks

- Confirming an alert on page 564

Confirming an alert

You can list, review, and delete alerts.

Procedure

1. On the **Tasks & Alerts** tab, click **All Alerts**.
2. Click on an alert to review the alert details and take appropriate action.

Result

The alert is resolved.

Tip: When the number of alerts for a storage system or file server node exceeds 10,000 alerts, memory might become insufficient. We recommend that you manually delete alerts that are no longer necessary. However, for the NAS Platform family, if the number of alerts for a single cluster exceeds 10,000, the alerts for that cluster are automatically deleted starting from the oldest alert.

As alerts are confirmed as resolved, select the resolved alerts, and click Delete Alerts.

Related concepts

- About alerts on page 563

A

Managing storage resources that have virtual IDs (VSP and HUS VM)

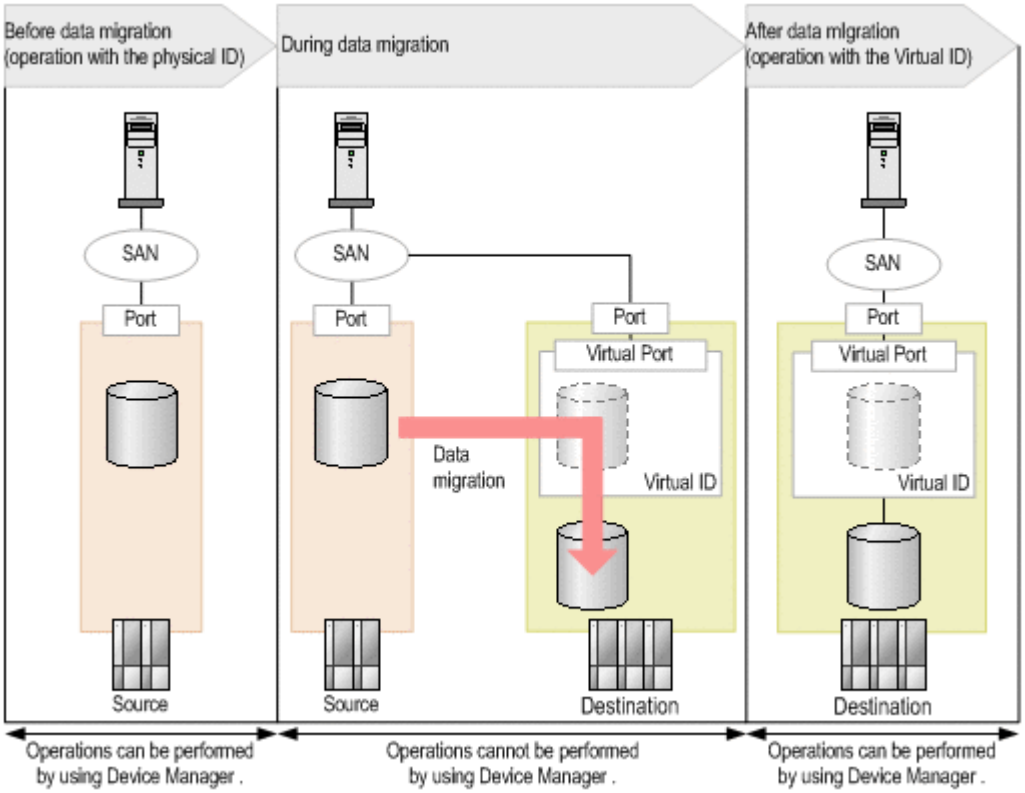
This module describes specific functionality available for virtualized volume resources for which virtual IDs are assigned.

- About managing storage resources that have virtual IDs
- Displaying virtual ID information
- Allocating volumes to hosts with virtual IDs

About managing storage resources that have virtual IDs

When replacing an old storage system with a Virtual Storage Platform or HUS VM, you can use virtual IDs to migrate host resources (volumes) without stopping the host. Virtual IDs provide information about the old storage system, and are configured on the migration-destination resource. By using Device Manager, you can check virtual IDs after data migration and perform operations on a resource that has been assigned virtual IDs. Even if you can view the virtual ID of a resource, until data migration to the target resource is complete, you cannot perform operations on the resource.

During the three stages of data migration, accessing the resource data changes from using the physical ID to the virtual ID, as the following graphic shows:



- **Before data migration (using physical IDs)**
  Before data migration, a host uses physical IDs to perform operations on the migration-source resource. Operations can be performed on the resource by using Device Manager.

- **During data migration**
  The path from the host to the storage system is changed from the migration source to the migration destination. Configure virtual IDs on the migration-destination resource, and then migrate the data from the migration-source resource. On the migration-destination storage system,

while data migration is in progress, operations cannot be performed on the resource by using Device Manager.

- **After data migration (using virtual IDs)**
  After all of the target resources are migrated, operations can be performed on a resource by using Device Manager. After data migration, a host uses virtual IDs to access resources and perform operations on the migration-destination storage system.

Items such as virtual LDEV ID and virtual ports can be used to verify virtual ID information.

Check the correspondence between the virtual ID information received from the host administrator and the physical ID information to identify the target resource.

You can perform any of the following tasks on resources that have virtual IDs:

- Expand DP volumes
- Reclaim zero pages of DP volumes
- Edit tier relocation of HDT volumes
- Edit tiering policy of HDT volumes
- Manage data placement profiles of HDT volumes
- Allocate volumes[1]
- Allocate like volumes[1]
- Edit LUN paths
- Unallocate volumes[2]
- Define copy pairs[3]
- Change the status of a copy pair[3]
- Migrate data by using Tiered Storage Manager
- Create logical group
- Create tier
- Edit volume labels
- Refresh volume labels

Notes:

1. For volumes for which no virtual ID is specified, you can specify new virtual IDs, and then allocate the volumes to host groups for which virtual IDs have already been specified.
2. You can also delete virtual IDs at the same time that volumes are unallocated.
3. To perform operations on a copy pair that contains volumes for which virtual IDs have been specified, you must update in advance the information (physical ID information) of the post-migration storage system with the configuration definition file of the copy pair. For details about the contents of configuration definition files, as well as the system configuration required to perform operations on copy pairs that contain

volumes for which virtual IDs have been specified, see the *Hitachi Command Suite Administrator Guide*.

Related concepts

- About removing hosts and related resources, or vCenters on page 79

Displaying virtual ID information

You can display virtual ID information in the volume list to identify the target resource from the virtual ID information. Virtual ID information of resources can be displayed per host, storage system, or logical group.

Procedure

1. On the **Resources** tab, select **Hosts**, **Storage Systems**, or **Logical Groups**.
2. Expand the tree and select the target resource to display the volume list.
3. Click the **Column Settings** button.
4. From the displayed dialog box, select check boxes to display virtual ID information, such as Virtual LDEV ID and Virtual Port.

Result

Virtual ID information is displayed in the list of volumes, enabling you to map the virtual ID with the ID (physical ID) information assigned by the destination storage system.

Tip: After checking the virtual ID information for the volume, we recommend editing the volume label so that the volumes can be identified from the virtual ID information, thus enabling you to identify the target volumes from windows that do not display virtual ID information.

Virtual ID information can also be displayed in the tables for Allocate volumes, Allocate like volumes, Unallocate volumes, or Edit LUN Paths dialog boxes. If virtual IDs are set but you cannot set columns in the tables in these dialog boxes, virtual ID information will be automatically displayed.

Allocating volumes to hosts with virtual IDs

You can allocate volumes to hosts that belong to resource groups being used for data migration and that are using virtual IDs. You can also allocate volumes for which no virtual IDs are specified. In this case, such volumes are assigned new virtual IDs. A resource group used for a data migration that uses virtual IDs is the resource group specified for the storage system when migrating data, not a resource group specified in Hitachi Command Suite.

Before you begin

- Identify the name of the target host.
- Identify the volumes to be allocated, if manually selecting such volumes.
- Identify the necessary volume capacity and size, if allocating volumes by specifying conditions.
- Verify that volumes to be allocated meet one of the following conditions:
  - The volumes belong to resource groups used for a data migration that uses virtual IDs, and the resource group is the same resource group to which the host group belongs.
  - The resource group ID of the storage system is 0 (it belongs to meta_resource). Note that meta_resource is a resource group specially managed on the storage system. It contains resources that are not a part of individual resource groups.

Procedure

1. On the **Resources** tab, select **Hosts**.
2. Expand the tree and select an operating system.
3. Select one or more hosts, and click **Allocate Volumes**.
4. Allocate volumes with existing virtual ID information, or volumes requiring virtual ID assignments as follows:
   a. To allocate volumes that belong to resource groups being used for a data migration that uses virtual IDs, reference the virtual ID information and manually select volumes, and then click **Show Plan** to register this as a task.
   b. To assign new virtual IDs to volumes that belong to meta_resource, and then allocate those volumes to a host, specify volume conditions or manually select volumes whose storage system resource group ID is 0 (meta_resource), and then click **Show Plan** to register this as a task. Use **Column Settings** to display storage system resource group IDs.
5. Confirm that the information in the plan summary is correct. If changes are required, click **Back**.
6. (Optional) Update the task name and provide a description.
7. (Optional) Expand **Schedule** to specify the task schedule. You can schedule the task to run immediately or later. The default setting is **Now**. If scheduled for **Now**, select **View task status** to monitor the task after it is submitted.
8. Click **Submit**. If the task is scheduled to run immediately, the task begins.
9. You can check the progress and the result of the task on the **Tasks & Alerts** tab. Click on the task name to view details of the task.

   Using the **Resources** tab, select the hosts to verify that the designated volumes are allocated. Volumes for which no virtual ID is specified are assigned a new virtual ID. When allocating volumes whose storage system resource group ID is 0, the volumes are registered in the same resource group as the host group.

Result

The designated volumes are allocated on the selected hosts.



Tip: When allocating like volumes, you can also allocate volumes for which no virtual IDs are specified, as described in this procedure.



Note: When unallocating volumes, you can also delete virtual IDs at the same time. To delete virtual IDs, specify conditions in Advanced Options > Virtual ID Information.

B

Linking related products

This module describes the link and launch feature that is used to access other related Hitachi Command Suite products.

- Launching other Hitachi Command Suite products
- Linking with Compute Systems Manager

Launching other Hitachi Command Suite products

Other integrated Hitachi Command Suite products for which you have a license installed and registered may be used via the link-and-launch feature.

About launching other Hitachi Command Suite products

By starting related products from HCS, you can configure more detailed settings or verify the whole image of the storage resource.

An integrated single sign-on is used for all HCS products for which you have an installed and registered license. Related HCS products must be installed on the same system as HCS to use the link-and-launch feature.

Link-and-launch enables you to easily move across storage management software for a full view of storage resources. HCS provides link-and-launch integration with the following software applications:

- Hitachi Tuning Manager
- Hitachi Global Link Manager
- Hitachi Replication Manager
- Storage Services Manager
- Protection Manager
- Hitachi Compute Systems Manager
- Hitachi Automation Director
- File server/NAS module management software
  You can start the file server/NAS module management software from Device Manager, which allows you to view detailed information about file servers/NAS modules and create file systems when allocating volumes.
  - For NAS Platform family, the file server management software is System Management Unit (SMU)/NAS Manager
  - For NAS Platform F and Data Ingestor the file server management software is File Services Manager (HFSM)
- Element Manager
  For storage systems other than VSP Fx00 models, VSP Gx00 models and VSP G1000, VSP G1500, and VSP F1500, you can start Element Manager from Device Manager, which allows you to view storage system information or remotely operate the storage system.

Note: If you have problems starting HCS products or Element Manager, follow the instructions in the displayed message, check the network status, and then try again, or report to the server administrator which product could not be activated.

Related tasks

- Starting related products on page 573

- Starting related products from a list of hosts on page 573
- Starting Element Manager on page 575
- Starting System Management Unit (SMU)/NAS Manager on page 576
- Starting File Services Manager (HFSM) on page 577

Starting related products

Before you can start related products, you must register their licenses.

Procedure

1. From the **Tools** menu, select the product you want to start.
2. Navigate to the newly launched browser to access the product.

Related concepts

- About launching other Hitachi Command Suite products on page 572

Related tasks

- Starting related products from a list of hosts on page 573

Starting related products from a list of hosts

You must first register the licenses of all products you want to start.

Protection Manager, Protection Manager Console, and a Device Manager agent must be installed on each host managed by Device Manager. For more information, see the *Protection Manager Console User's Guide*.

Procedure

1. From the **Resources** tab, select **Hosts**.
2. From the **Actions** menu located in the corner of the application pane, select **Protection Manager**.
3. Navigate to the newly launched browser to access the product.

Related concepts

- About launching other Hitachi Command Suite products on page 572

Related tasks

- Starting related products on page 573

Notes on starting Element Manager

When starting the storage system software (Element Manager) from Hitachi Command Suite, keep the following in mind.

For HUS 100, AMS 2000, SMS, and AMS/WMS:

- When Element Manager is left inactive for a certain period of time, auto log off is executed, and the user cannot use the software. If auto log off is

executed when password protection is enabled, the user remains logged in to the storage system, and cannot perform a refresh or create volumes on the storage system from HCS. In this case, restart Element Manager, and then manually close the window.
For details about auto log off, see the manuals for each storage system.

For HUS 100, AMS 2000, and SMS:

- Element Manager can be started from one management client for each storage system. Do not start Element Manager from multiple management clients at the same time.

For HUS 100 and AMS 2000:

- When Account Authentication is used, executing the following operations while Element Manager is being started causes the user to remain logged in to the storage system, preventing operations on storage systems from Element Manager and HCS until the session times out:
  - Log out from HCS
  - Stop the services of Hitachi Command Suite common components and Hitachi Command Suite products by using the **hcmdssrv** command.
  - Close Element Manager with nodes (such as components and groups) below the storage system name selected on the navigation tree of Element Manager.
    Closing Element Manager with the storage system name (route node) selected on the navigation tree causes the user to be logged out.

  For details about the **hcmdssrv** command, see the *Hitachi Command Suite Administrator Guide*.
  For details about session timeout times, see the manuals for each storage system.

For AMS and WMS:

- To use Firefox, you need to disable the setting that blocks non-SSL active content that coexists on SSL pages. In the address bar of Firefox, enter `about:config` to open the configuration window, and then change the value of `security.mixed_content.block_active_content` to `false`.

For SMI-S enabled storage systems:

- The management tool sometimes cannot start with an SMI-S enabled storage system because the URL of management tool is incorrect. In this case, you need to perform either of the following operations:
  - Use the **AddURLLink** command of Device Manager CLI to register in the Device Manager server the URL of the management tool of the SMI-S enabled storage system. For details about the **AddURLLink** command, see the *Hitachi Command Suite CLI Reference Guide*.
  - Review the settings of the SMI-S provider, and then refresh the SMI-S enabled storage system from Device Manager.

For Virtual Storage Platform:

- If SSL or TLS communication is used between Virtual Storage Platform and the HCS management server, the startup of Element Manager might fail due to a communication error. In that case, check the server certificate of the storage system. If MD5 is used as a signature algorithm, recreate the server certificate by using SHA (SHA-256 is recommended).

Related tasks

- Starting Element Manager on page 575

Starting Element Manager

Element Manager is a menu selection for launching native storage system management tools, which allows you to remotely view, configure, and operate storage systems, including SMI-S enabled storage.

Element Manager is available for enterprise storage systems, and specific mid-range storage systems, as soon as the storage system is registered with an IP address and user login information.

This includes Virtual Storage Platform, Universal Storage Platform V/VM, and mid-range systems such as HUS 100, AMS 2000, and SMS.

For example, after Virtual Storage Platform discover and registration completes, you are able to launch Storage Navigator using the Element Manager menu selection.

If Element Manager does not display in the Actions menu, the mid-range storage system may require environmental settings. See the following prerequisites.

Before you begin

Configure the environment for starting Element Manager.

For details, see Element Manager in the *Hitachi Command Suite Administrator Guide*.

Procedure

1. On the **Resources** tab, select **Storage Systems**.
2. Select a storage system.
3. From the **Actions** menu in the application pane, select **Element Manager**.

Result

Confirm the storage system management tool is started and available for use.

Related tasks

- Starting System Management Unit (SMU)/NAS Manager on page 576
- Starting File Services Manager (HFSM) on page 577

Related references

- Notes on starting Element Manager on page 573

Starting System Management Unit (SMU)/NAS Manager

In HCS, you can start the SMU/NAS Manager software to perform NAS Platform management tasks.

Before you begin

Register the NAS Platform family file server or VSP G400, G600, G800 or VSP Fx00 models with NAS modules in Device Manager.

Procedure

1. From the **Resources** tab, select **File Servers > All File Servers**. For NAS modules, select **Storage Systems > All Storage Systems**.
2. In the Server/Cluster list, select the NAS Platform cluster or file server, and click **NAS Manager**.

Result

The browser opens running SMU/NAS Manager.

Note:

- If the IPv6 address of the managing server is registered in SMU/NAS Manager version 12.x, SMU/NAS Manager cannot be started. In such a case, start SMU/NAS Manager directly instead of starting it from Device Manager.
- Except for NAS Manager in the NAS modules of VSP G400, G600, G800, or VSP Fx00 models, use the account created by using NAS Manager. For NAS Manager in the NAS modules of VSP G400, G600, G800, or VSP Fx00 models, use the account created by using Device Manager - Storage Navigator.
- For NAS Manager in the NAS modules of VSP G400, G600, G800, or VSP Fx00 models, if you use an account that was created in a firmware version earlier than 83-03-2X-XX/XX and the firmware was updated to version 83-03-2X-XX/XX or later, then change the password. Otherwise, you will be unable to log in to NAS Manager.

Related tasks

- Starting File Services Manager (HFSM) on page 577
- Starting Element Manager on page 575

Starting File Services Manager (HFSM)

In HCS, you can start File Services Manager software to perform NAS Platform F or Data Ingestor (HDI) management tasks.

Before you begin

NAS Platform F or Data Ingestor file server in Device Manager.

Procedure

1. From the **Resources** tab, select **File Servers**, then **All File Servers**.
2. Open the tree and select a NAS Platform F or Data Ingestor node, and from the **Actions** menu in the corner of the application pane, select **File Server Manager**.

Result

The browser opens running File Services Manager.

Related tasks

- Starting System Management Unit (SMU)/NAS Manager on page 576
- Starting Element Manager on page 575

Linking with Compute Systems Manager

This module describes how to link to Compute Systems Manager.

About linking with Compute Systems Manager

If the host management product Compute Systems Manager (HCSM) is installed on the Device Manager server and linked, information about hosts being managed by both products will be synchronized automatically when using Host Data Collector.

When Device Manager and Compute Systems Manager are linked, hosts are registered or updated simultaneously between the two products. For this reason, Compute Systems Manager hosts can automatically be registered or updated as hosts of Device Manager. Similarly, hosts that have been registered or updated in Device Manager can automatically be registered or updated in Compute Systems Manager. Therefore, it is no longer necessary for both the storage and server administrators to separately register or update the same hosts.

To synchronize host information, the requirements are:

- Hosts must have either a WWN or an iSCSI name.
  Among the Compute Systems Manager hosts, information about those that have a WWN or an iSCSI name is sent to Device Manager.Compute

Systems Manager does not obtain iSCSI names when hosts are registered, but Device Manager hosts that have iSCSI names are also registered in Compute Systems Manager.

- Hosts must be normal hosts or virtual machines whose information was obtained by using Host Data Collector.
  Among Device Manager hosts, information about hosts that were registered using Host Data Collector is sent to Compute Systems Manager. Information about Compute Systems Manager hosts is also sent to Device Manager through Host Data Collector. Because Device Manager does not support the `su` command, if you plan to access Linux hosts by using the `su` command from Compute Systems Manager, then in advance, ensure that settings are configured for the `sudo` command from Device Manager to ensure that host information is properly synchronized.
  Hosts that were registered by using a method other than Host Data Collector (such as a Device Manager agent or a host scan) and virtualization servers are not synchronized. For details about setting up Host Data Collector and setting the `sudo` command, see the *Hitachi Command Suite Administrator Guide*.

Note that even when Device Manager and Compute Systems Manager are linked, host deletions are not synchronized. Therefore, deleting a host from Device Manager does not affect Compute Systems Manager. Compute Systems Manager hosts that have been deleted from Device Manager can be manually re-synchronized, if necessary.

You can also start and perform operations in Compute Systems Manager from Device Manager by using the single sign-on functionality.

Related concepts

- Methods for registering hosts on page 66

Related tasks

- Resynchronizing hosts of Compute Systems Manager on page 578

Resynchronizing hosts of Compute Systems Manager

Users can re-synchronize Compute Systems Manager hosts that have been deleted from Device Manager.

Procedure

1. In the **Administration** tab, select the **Hosts** tab, and click **Add Hosts**.
2. Select **Re-sync Hosts with CSM**, and then submit the task.
   The task is registered in the **Data Collection Tasks** tab.
3. Check the task status in the **Data Collection Tasks** tab.

Result

If the task completed, a host is added to the Hosts tab.

Related concepts

- About linking with Compute Systems Manager on page 577

Glossary

A

allocated volume

A logical device (LDEV) for which one or more host LUN paths are defined.

C

C/T

See consistency time.

C/T delta

The consistency time difference, or write-delay time, of data transfer between the primary and secondary storage systems.

cache

A set of RAM (Random Access Memory) modules used to store data temporarily.

cache logical partition (CLPR)

Virtual cache memory that is set up to be allocated to hosts that are in contention for cache memory. CLPRs can be used to segment storage system cache that is assigned to parity groups.

capacity

The amount of data storage space available on a physical storage device, generally measured in bytes (MB, GB, TB, and so on).

CCI

See Command Control Interface.

CLI

command line interface

CLPR

See cache logical partition.

Command Control Interface (CCI)

Software used to control volume replication functionality (such as TrueCopy or ShadowImage) by means of commands issued from a host to a storage system. A command device must be set up in the storage system to enable the storage system to receive commands from CCI.

In an open system, Replication Manager uses the CCI configuration definition files to modify copy pair configurations and to acquire configuration information. Copy pair modification processing, such as splitting and resynchronizing copy pairs, is executed on the storage system via CCI.

concatenated parity group

The concatenation of two or more parity groups into one group. Using a concatenated parity group reduces the time that is required to access data (especially sequential data).

consistency time (C/T)

A replication policy or threshold that indicates the amount of time that a replication target (volume, journal group, or extended consistency group, for example) is allowed to lag behind replication of the master, or source, volume.

control unit (CU)

Created in an enterprise-class storage system. Also called a CU image. The LDEVs created in a storage system are connected to a single CU, and a number is assigned to each CU for identifying its LDEVs. Therefore, volumes (LDEVs) in a storage system are specified by the CU number (CU#) and LDEV number.

copy pair

A primary and secondary volume pair linked by the volume replication functionality of a storage system. The primary volume contains original data, and the secondary volume contains the copy of the original.

Copy operations can be synchronous or asynchronous, and the volumes of the copy pair can be located in the same storage system (local copy) or in different storage systems (remote copy).

CSV

comma-separated values

CU

See control unit.

D

DB

database

deduplication system data volume

The volume used to manage data duplication in a pool. The deduplication system data volume is created when you enable deduplication on a pool.

device (dev or DEV)

A physical or logical unit with a specific function.

DEVN

Device number that is assigned to each logical address when using an LDEV on a mainframe host.

DHCP

Dynamic Host Configuration Protocol

discovery

A process that finds and identifies network objects. For example, discovery may find and identify all hosts within a specified IP address range.

DKP

Disk processor

DP

Hitachi Dynamic Provisioning. See *Dynamic Provisioning*.

DP pool

The area where DP pool volumes (actual volumes) are registered. When a DP volume (virtual volume) receives a write operation from a host, that data is stored in a DP pool volume.

When Dynamic Provisioning and Dynamic Tiering must be differentiated, this document uses the terms HDP pool and HDT pool.

DP pool volume

An actual volume that is one of the volumes making up a DP pool.

When Dynamic Provisioning and Dynamic Tiering need to be distinguished, this manual uses the terms HDP pool volume or HDT pool volume.

DP volume

A virtual volume that is created from a Dynamic Provisioning (DP) pool (that is, it is associated with a DP pool).

When Dynamic Provisioning and Dynamic Tiering must be differentiated, this document uses the terms HDP volume and HDT volume.

drive

A physical data storage device that can be either a hard disk drive (HDD) or a flash (solid-state) drive.

DSD volume

Used in Hitachi Command Suite to identify a *deduplication system data volume*.

E

EVS

Enterprise Virtual Server

external path

A path from a storage port of a storage system to a volume on a connected external storage system.

external volume

A logical volume whose data resides on drives that are physically located in an externally connected storage system.

F

FC

Fibre Channel

FCoE

Fibre Channel over Ethernet. An encapsulation of Fibre Channel frames over Ethernet networks. This allows Fibre Channel to use 10-gigabit Ethernet networks (or higher speeds) while preserving the Fibre Channel protocol.

flash module drive (FMD)

A storage device, which is developed by Hitachi, that uses flash memory.

FMD

See flash module drive.

G

global-active device

A storage system feature that provides high-availability copy pairs to support simultaneous, uninterrupted host and cluster active-active I/O within and across geographically disparate data centers. This feature protects against storage system and site failures by providing continuous access to data.

GUI

graphical user interface

H

HBA

See host bus adapter.

HDT

See Hitachi Dynamic Tiering.

Hitachi Dynamic Provisioning (HDP)

Functionality that allocates virtual volumes to a host and uses the physical capacity that is necessary according to the data write request.

Hitachi Dynamic Tiering (HDT)

Functionality that is used with Hitachi Dynamic Provisioning that places data in a hardware tier according to the I/O load. For example, a data area that has a high I/O load is placed in a high-speed hardware tier, and a data area that has a low I/O load is placed in a low-speed hardware tier.

host bus adapter (HBA)

One or more dedicated adapter cards that are installed in a host, have unique WWN addresses, and provide Fibre Channel I/O connectivity to storage systems, typically through Fibre Channel switches. Unlike general-purpose Ethernet adapters, which handle a multitude of network protocols, host bus adapters are dedicated to high-speed block transfers for optimized I/O performance.

host group

Custom grouping of hosts that segregates hosts in a meaningful way, for example, a group of hosts that is segregated by operating system. A host group can be shared with another virtual port or another physical port for alternate path support.

HTML

Hypertext Markup Language

HTTP

Hypertext Transfer Protocol

HTTPS

Hypertext Transfer Protocol Secure

I

I/O

input/output

internal volume

A logical volume whose data resides on drives that are physically located within the storage system.

Internet protocol (IP)

The protocol that governs the breakup of data messages into packets (units of data), the routing scheme for transmitting them, and the reassembly of the packets into the original data messages at the destination. Most networks combine IP with a higher-level protocol called Transmission Control Protocol (TCP), which establishes a virtual connection between a source and a destination.

IOPS

I/Os per second

IP

See Internet protocol.

iSCSI

Internet Small Computer Systems Interface

J

JRE

Java Runtime Environment

JVM

Java Virtual Machine

JWS

Java Web Start

L

LAN

See local area network.

LDAP

Lightweight Directory Access Protocol

LDEV

See logical device.

LDKC

Logical disk controller

local area network (LAN)

A computer network that spans a relatively small geographic area, such as a single building or group of buildings.

logical device (LDEV)

A volume created in a storage system. See also LU.

logical group

A user-defined collection of managed resources (hosts and volumes) that are grouped according to business operations, geographic locations, or other organizational divisions. Logical groups can be *public* or *private*:

- Public logical groups are accessible by any HCS user.
- Private logical groups are accessible only by HCS users who belong to user groups that are associated with the logical group.

logical unit (LU)

A volume, or LDEV, created in an open storage system, or configured for use by an open-systems host, for example, OPEN-V.

logical unit number (LUN)

A unique management number that identifies a logical unit (LU) in a storage system. A logical unit can be an end user, a file, a disk drive, a port, a host group that is assigned to a port, an application, or virtual partitions (or volumes) of a RAID set.

Logical unit numbers (LUNs) are used in SCSI protocols to differentiate disk drives in a common SCSI target device, such as a storage system. An open-systems host uses a LUN to access a particular LU.

LU

See logical unit.

LUN

See logical unit number.

LUSE

LUN expansion; LU size expansion

M

M-JNL

See master journal.

main control unit (MCU)

A storage system at a primary, or main, site that contains primary volumes of remote replication pairs. The main control unit (MCU) is configured to send remote I/O instructions to one or more storage systems at the secondary, or remote, site, called remote control units (RCUs). RCUs contain the secondary volumes of the remote replication pairs. See also remote control unit (RCU).

management client

A computer used to operate a graphical user interface client or a command-line interface client.

master journal (M-JNL)

The primary, or main, journal volume. A master journal holds differential data on the primary replication system until the data is copied to the restore journal (R-JNL) on the secondary system. See also restore journal.

MCU

See main control unit.

N

NAS

Network attached storage

NIC

Network interface card

P

P-VOL

See primary volume.

pair status

Indicates the condition of a copy pair. A pair must have a specific status for specific operations. When a pair operation completes, the status of the pair changes to a different status determined by the type of operation.

path

A path from a storage system volume to a host group.

In this manual, the term "path" may mean a path, external path, or LUN path without making distinctions among them.

pool volume (pool-VOL)

A logical volume that is reserved for storing Copy-on-Write Snapshot data or Dynamic Provisioning write data.

primary volume (P-VOL)

In a volume pair, the source volume that is copied to another volume using the volume replication functionality of a storage system. The data on the P-VOL is duplicated synchronously or asynchronously on the secondary volume (S-VOL).

properties file

A file that defines aspects of the operating environment. The operating environment can be modified by changing the appropriate properties file.

R

R-JNL

See restore journal.

RACF

See Resource Access Control Function.

RADIUS

Remote Authentication Dial In User Service

RAID

redundant array of independent disks

A collection of two or more disk drives that presents the image of a single logical disk drive to the system. Part of the physical storage capacity is used to store redundant information about user data stored on the remainder of the storage capacity. In the event of a single device failure, the data can be read or regenerated from the other disk drives.

RAID employs the technique of disk striping, which involves partitioning each drive's storage space into units ranging from a sector (512 bytes) up to several megabytes. The stripes of all the disks are interleaved and addressed in order.

RAID level

The type of RAID implementation. RAID levels include RAID 0, RAID 1, RAID 2, RAID 3, RAID 4, RAID 5 and RAID 6.

RCU

See *remote control unit (RCU)*.

recovery point objective (RPO)

In business continuity planning, the maximum tolerable period in which data can be lost from an IT service due to a major incident.

refresh

To update the database using the most recent information.

remote control unit (RCU)

A storage system at a secondary, or remote, site that is configured to receive remote I/O instructions from one or more storage systems at the primary, or main, site. See also main control unit.

Resource Access Control Function (RACF)

Functionality for controlling user authentication and resource access on the mainframe host.

resource group

A collection of resources that are grouped by one or more system resource types.

resource pool

A type of resource group to which resources of a virtual storage machine in a VSP G series or VSP F series storage system belong, if those resources have not been added to another resource group. There are two types of resource pools: the resource pool on the default virtual storage machine and the resource pool that is created automatically for each virtual storage machine that you create on the storage system.

restore journal (R-JNL)

The secondary, or remote, journal volume. A restore journal holds differential data on the secondary replication system until the data is copied to the secondary volume (S-VOL). See also master journal (M-JNL).

role

Permissions that are assigned to users in a user group to control access to resources in a resource group. Resource groups can be assigned to different user groups with different roles.

RPO

See recovery point objective.

S

S-VOL

See secondary volume.

SAN

See storage area network.

SAS

See Serial Attached SCSI.

SATA

See Serial Advanced Technology Attachment.

SCSI

Small Computer System Interface. A standard that defines I/O buses primarily intended for connecting storage systems and devices to hosts through host bus adapters.

secondary volume (S-VOL)

After a backup, the volume in a copy pair that is the copy of the original data on the primary volume (P-VOL). Recurring differential data updates keep the data in the S-VOL consistent with the data in the P-VOL.

Secure Sockets Layer (SSL)

A common protocol for managing the security of message transmission over the Internet.

Two SSL-enabled peers use their private and public keys to establish a secure communication session, with each peer encrypting transmitted data with a randomly generated and agreed-upon symmetric key.

Serial Advanced Technology Attachment (SATA)

A computer bus technology primarily designed for the transfer of data to and from hard disks and optical drives. SATA is the evolution of the legacy Advanced Technology Attachment (ATA) interface from a parallel bus to serial connection architecture.

Serial Attached SCSI (SAS)

A replacement for Fibre Channel drives in high-performance applications. See also SCSI.

SLPR

Storage Local Partition

Small Computer System Interface

See SCSI.

SMU

System Management Unit

SNMP

Simple Network Management Protocol

SSD

Solid-state drive

storage area network (SAN)

A network of shared storage devices that contain disks for storing data.

storage pool

A collection of system drives or the logical container of a file system. Storage pools are created in the NAS Platform family of products. Storage pools are also called *spans*.

system drive

The basic (logical) storage element that is managed by the NAS Platform family of products. A system drive is equivalent to a storage system volume. The load on volumes is distributed among system drive groups, which are created by users.

T

tiered storage

A layered structure of performance levels, or tiers, that matches data access requirements with the appropriate performance tiers.

U

unallocated volume

A volume (LDEV) for which no host paths are assigned.

URL

Uniform Resource Locator

user group

A collection of users who have access to the same resources and have the same permissions for those resources. Permissions for users are determined by the user groups to which they belong. Users and resource groups can be assigned to multiple user groups.

V

virtual storage machine

A virtual storage system that you create on a VSP G series or VSP F series storage system that allows multiple resource groups to be treated as a single device.

virtual volume (V-VOL)

A logical volume in a storage system that has no physical storage space. Hitachi Thin Image uses V-VOLs as secondary volumes of copy pairs. In Hitachi Dynamic Provisioning, V-VOLs are referred to as DP-VOLs.

VOLSER

The label of a volume assigned by the mainframe host.

volume (vol or VOL)

A name for the logical device (LDEV), or logical unit (LU), or concatenated LDEVs, that are created in a storage system that have been defined to one or more hosts as a single data storage unit.

W

web client

An application that is used on a client machine to access a server on which management software is installed. A web client contains two parts: dynamic web pages and the web browser.

WWN nickname

World wide name nickname. A name that is set for the WWN of an HBA that identifies which HBA to operate on. The WWN nickname easily identifies accessible HBAs that are registered in a host group. You can display a list of WWN nicknames to confirm target HBAs for a host while you edit LUN paths during or after volume allocation or when you replace an HBA.

Index

A

accelerated compression, enabling 149
access attribute
  changing to read/write 186
access control 88, 100
  users 99
accounts
  creating 91
  editing 93, 94
  lock status 95
  locking (automatic) 45
  managing (overview) 88
  strong passwords 44
adding redundancy
  global-active device 356
agents
  downloading 46
alerts
  confirming 564
  deleting 564
  description 563
allocating existing volumes
  global-active device 353
allocating new pairs
  global-active device 350
ALUs
  determining volume attributes for 418
analytics tab
  display differences 550
authentication
  external, groups 97
  external, users 96

B

backing up a local copy 317
bottleneck
  detection 530, 546
  MP blade, unit 532
  overview 526
browser
  clearing cache 41
  configuring (overview) 34

C

capability profiles
  use of default 408
checking status
  global-active device 358
CLI
  downloading 46
CLPR
  creating 439
  deleting 443
  editing 441
  migrating 440, 443
clustered-host storage
  (overview) 223
  creating 224
command devices
  global-active device 348
Compute System Manager 578
  linking 577
copy pairs
  defining 282
  managing (overview) 281, 288
  operations 281, 288
  status, changing 283, 284
  status, deleting 285
copy topology
  changing name 306
  creating 303, 305
  deleting 314
  expanding 310, 311
  expanding by changing type 308
  shrinking 312, 313
  shrinking by changing type 309
copy topology type
  about changing 307
copy-back mode, editing 148
creating
  storage containers 418
  VVols 416
creating a report 501
creating file systems, NAS 264
criteria search
  dialog box 494
CSV file 499
  file server 498

hosts 497, 500
logical groups 499
types 497
custom
roles 111

D

dashboard
accessing 521
customizing 522
description 519
data collection tasks
restarting 515
status 514
Data Collector
downloading 46
data migration
dialog box 462
data mobility
overview 444
data optimization
logical group 449
pools 449
data transfer speed
Fibre Channel ports 235
deleting
protocol endpoints 425
storage containers 423
deleting a report 502
disaster recovery (testing) 318, 319
discontinue
global-active device 365
download software 46

E

Editing port settings 233
element manager
notes 573
starting 575
Element Manager
starting products 573
EVS, updating HCS linkage 86

F

failures
global-active device 363
Fibre Channel ports
configuring 237
data transfer speed 235
file server
adding volumes 206
export CSV 498
unallocating volumes 228
file server, name change 85
file servers
registering 84
file servers, registering 81
file storage, removing VVol environment 430
file storage, VVol configuration for 426
file systems
creating, NAS Platform 246
Firefox
security 38
functions
roles, permissions 114

G

global-active device
adding redundancy 356
allocating existing volumes pairs 353
allocating new pairs 350
checking status 358
command devices 348
discontinue 365
failures 363
host path failure 364
monitoring pairs 359
overview 334
prerequisites 342
protecting pairs 334
quorum disk 345
replication paths 343
selecting storage 342
setup window 340
unallocating 360
unallocating volume 362
virtual storage machine 349
Google Chrome
setting options for 39
guidelines, for VVol resources 414

H

HBA
add 436
optimize 435
remove 244, 245
replace, exchange 243
HCS
components 22
configuration (overview) 34
features 23
features, new 24
login, licensing 41
navigating the interface 29
process flow 28
system configuration 25
HDI, create file systems HFSM 3.1 277
health check
edit schedule 542
exporting 542
history 543

scheduling 540
health check reports, notes on 528
Help
navigating 31
high temperature mode
enabling 438
high temperature mode, about 437
HNAS F or HDI file system
deleting 279
HNAS F/HDI
file systems/shares 273
HNAS F/HDI, create file systems HFSM 3.2+ 276
HNAS F/HDI, file system settings HFSM 3.1+ 278
HNAS F/HDI, HFSM workflow 275
HNAS F/HDI, launching HFSM 577
host path failure
global-active device 364
host workflow
detecting, merging 74
hosts
change manually registered 76
device manger updates 77
exporting 497, 500
HDC updates 75
registering, about 66
registering, HDC 69
registering, host scan 71
registering, merge host 72
registering, methods 66
registering, WWN/iSCSI 70
removing 80
removing, about 79
settings, information 73
starting products 573
vCenter refresh 78
viewing 525
Hosts view (overview) 524

I

IE browser
pop-up blocking 36
Internet Explorer
security 37

J

Java Web Start
clearing cache 41
log output 136
multiple JREs 40
Java Web Start 6.0/7.0 39

L

launching
optional software 572
launching other products 572
license
management 46
registering 47
license keys
types 48
viewing information 52
License keys
disabling 50
licenses
checking 48
licenses, viewing 52
log output 136
logging in 42
logging on 41
logical groups
creating 477
defined 472
deleting 480
editing 479
exporting 499
operations, performing 479
reports, viewing 478
LUN paths
editing host group information 242
editing paths in host clusters 240
editing WWN nicknames 240, 242
managing WWN nicknames 238
LUN security
enabling on ports 231
LUN security on ports
disabling 232
LUSE
expansion 157, 158
releasing 159

M

managing (overview) 100
metrics 538, 539, 545, 546, 556
metrics, health check 544
metrics, performance wizard 537
migration
about projects and jobs 370
allocating target S-VOLs 390
changing host I/O paths 383
changing job descriptions 399
changing job names 399
changing status of failed jobs 396
creating a project 376
creating copy pairs 391
creating jobs 378
delete source volumes 401
deleting a job 399
deleting job volumes 400
deleting projects 400
editing projects 398
job status 394
migrating config files 390
migrating volume data 386

migration job workflow 380
optimizing target storage 403
overview 368
pre-migration planning 369
preparing storage resources 375
preparing target S-VOLs 388
project status 393
removing source storage 402
resuming failed jobs 396
shred source volumes 401
unallocate source volumes 385
unvirtualize source volumes 387
user interface overview 371
virtualizing source volumes 382
zero page reclaim 402
modifying
storage containers 422
monitoring
overview 518
monitoring pairs
global-active device 359
MP Blade
viewing 524

N

name resolution
management server 35
NAS
configuring 254
NAS Platform file shares
deleting 272
NAS Platform file system
deleting 269
NAS Platform storage pool
deleting 264
NAS Platform, launching SMU/NAS Manager 576
NAS, add file shares 270
NAS, changing file system settings 267
NAS, changing storage pool name
11.3 or later 260
NAS, creating storage pools
earlier than v11.3 261
v11.3 or later 257
NAS, DP pools 253
NAS, edit file shares 271
NAS, environment settings 82
NAS, expand file system 266
NAS, expand storage pool
v11.3 or later 259
NAS, expanding storage pools
earlier than v11.3 262, 263
NAS, file system capacity 256
NAS, LUN paths 257
NAS, mount 268
NAS, parity groups pools 253
NAS, unmount 268
NAS, Workflow 250
NAS/HDI, file share concepts 279
navigating
Help 31

O

overview
global-active device 334
of SPBM 407
of VMware VVols 406

P

P-VOL Write Transfer Rate 562
pair management 281, 288
parity group
creation 142, 143
deleting 145
deletion 145
parity groups
create, dialog box 144
password
changing 94
path management 238, 488
performance
optimizing by setting data transfer speed for a Fibre Channel port 235
performance reporter
display configuration 548
display differences 550
permissions
by function 114
changing 95
policy
User ID, password 92
pools
deleting 184
threshold, modifying 184
creating 173, 183
creating, dialog box 176
expanding or converting 180
expanding volumes 185
overview 171
verifying 180
port scan and volume discovery 163
prerequisites
global-active device 342
profile
editing 93, 94
protecting pairs
global-active device 334
protocol endpoints
defining 417
deleting 425

Q

quorum disk
global-active device 345

R
reclaiming zero pages 186
removing
VMware VVol environment 423
removing a software application 51
removing, VMware VVol environment 424
replication
Continuous Access Journal wizard 559
UR wizard 559
analytics, Continuous Access Journal expert 560
analytics, UR expert 560
metric thresholds 555
performance 553
Replication, bandwidth 555
reports 561
replication configuration
alerts 314
backups and restores (local copy) 317
changing status 316
creating 293
disaster recovery test 318, 319
discontinuing 300
DP pools 295
journal groups 296
monitoring 315
pair management servers 297
quorum disk 298
remote paths 294
takeover and takeback 318
Thin Image pools 297
unsupported topology 312, 313, 317
virtual storage machine 299
workflow 292
Replication Manager 281
replication paths
global-active device 343
Replication tab 288
site management 302
Report Viewer window 501
reports
downloading 500
DP pools, ranking 447
logical groups, ranking 446
types 497
using 500
resource group configuration
required for VVols 414
resource groups
creating 104
deleting 106
editing 105
overview 102
prerequisites 104, 326
resource labels
editing 487
importing 488
overview 486
searching 487
restoring a local copy 317
resynchronizing hosts 578
roles
custom 111
user 110

S
search
exporting CSV file 499
security
about 44
account lock status 95
account locking 45
Firefox 38
Internet Explorer 37
strong passwords 44
warning banner 45
selecting storage
global-active device 342
setting
T10 PI mode on a port 234
setup window
global-active device 340
SLUs
determining volume attributes for 418
SMI-S 466, 467
software
downloading 46
software installation 49
spare drives
assigning 147
removing 148
SPBM
implementation examples 408
overview of 407
storage
allocation (overview) 141
monitoring (overview) 518
optimization, about 434
refreshing status 63
registering (overview) 57
searching (overview) 491
tier-based management 481
virtualization with pools 171
virtualized 140
virtualized tiers 187
storage containers
creating 418
deleting 423
expanding capacity 420
modifying 422
storage containers, creating for file storage VVols 427
storage containers, deleting for file storage VVols 430
storage containers, expanding the capacity for file storage VVols 429
storage containers, modify for file storage VVols 429

Storage Navigator
reports 500
storage performance by host
analyzing 534
storage resources
searching 492, 493
storage systems
editing 62, 136
logical groups 479
refreshing 64
registering 60
registering, prerequisites 57
removing 63
viewing 523
Storage view
about 523
supported functions
for VVol managment 413
system tasks
troubleshooting 513
status 511

T

T10 PI mode
setting on a port 234
takeover and takeback 318
task workflow
for VMware VVol environment 409
tasks
about 504
canceling (scheduled) 509
data retention 507
deleting 511
failed 509, 512
rescheduling 508
restarting failed or stopped 510
status 507
stopping (running) 508
summary, moving 509, 512
tiers
creating 481
deleting 483
expanding 482
performance 481
storage management 481
TLS settings
enabling 39

U

unallocating
global-active device 360
unallocating volume
global-active device 362
Universal Replicator
performance 552, 562
upgrade
clearing cache 41
UR Performance Analysis window
prerequisites 553
refresh 563
user
roles 110
user accounts
deleting 98
password, changing 94
passwords 94
user group
editing 135
user groups
creating 134
deleting 136
overview 107
users
access control 99
creating 91
lock status 95
locking (automatic) 45
managing (overview) 88
password, changing 94
permissions, changing 95
profile, editing 93, 94

V

VASA provider, changing for file storage VVols 428
viewing a report 500
virtual IDs
displaying 568
volume allocation 568
volume migration 566
virtual storage machine
allocating volumes 328
global-active device 349
virtual storage machines
about 322
creating 327
deleting 331
editing 330
virtualize volumes dialog box
volumes, virtualize 166
virtualizing tiers
overview 187
VMware VVols
overview of 406
removing environment for 423
VMware VVols, removing environment 424
volume
allocation 210, 212
volume attributes
determining for ALUs and SLUs 418
volumes
allocating 204, 205, 208, 209, 211, 224
allocating (overview) 155, 201
allocating, dialog box 213
allocating, prerequisites 202, 204

creating 151, 157
creating (overview) 149, 156
creating, dialog box 152
data migration 453
data migration, notes 452
data placement profile, applying 198
data placement profile, creating 195
data placement profile, editing 197
data placement profile, scheduling 199
data placement profile, updating 196
data placement, notes 194
migrating (overview) 451
migrating data in, prerequisites 200, 459, 460
pair management 281, 288
port scan 163
quick format 151
shredding 155
tier relocation, editing 192
tier relocation, monitoring 190
tier relocation, scheduling 191
tier relocation, starting 191
tiering policy, applying 193
tiering policy, customizing 193
unallocating 227
unallocating (overview) 226
unallocating, dialog box 229
unallocating, prerequisites 227
unvirtualizing 165
virtualizing (overview) 160
virtualizing discovered external storage system 162
virtualizing undiscovered external storage system 163
VVol resources, guidelines for 414
VVols
creating 416
related to block storage node 423
related to storage container node 423
supported functions for management 413

W

warning banner
setting 45
Web client
configuring (overview) 34
WWN
acquiring, conditions 68
nicknames 63, 213, 238
WWN nickname
editing 490
WWN nicknames
overview 488

Hitachi Data Systems

Corporate Headquarters
2845 Lafayette Street
Santa Clara, California 95050-2639
U.S.A.
www.hds.com

Regional Contact Information

Americas
+1 408 970 1000
info@hds.com

Europe, Middle East, and Africa
+44 (0) 1753 618000
info.emea@hds.com

Asia Pacific
+852 3189 7900
hds.marketing.apac@hds.com

Hitachi Data Systems

MK-90HC172-25